

---

# Towards Agnostic Feature-based Dynamic Pricing: Linear Policies vs Linear Valuation with Unknown Noise

---

Jiayu Xu

Computer Science Department,  
University of California, Santa Barbara

Yu-Xiang Wang

Computer Science Department,  
University of California, Santa Barbara

## Abstract

In feature-based dynamic pricing, a seller sets appropriate prices for a sequence of products (described by feature vectors) on the fly by learning from the binary outcomes of previous sales sessions (“Sold” if valuation  $\geq$  price, and “Not Sold” otherwise). Existing works either assume *noiseless* linear valuation or *precisely-known* noise distribution, which limits the applicability of those algorithms in practice when these assumptions are hard to verify. In this work, we study two more agnostic models: (a) a “linear policy” problem where we aim at competing with the best linear pricing policy while making no assumptions on the data, and (b) a “linear noisy valuation” problem where the random valuation is linear plus an unknown and assumption-free noise. For the former model, we show a  $\tilde{O}(d^{\frac{1}{3}}T^{\frac{2}{3}})$  minimax regret up to logarithmic factors. For the latter model, we present an algorithm that achieves an  $\tilde{O}(T^{\frac{3}{4}})$  regret, and improve the best-known lower bound from  $\Omega(T^{\frac{3}{5}})$  to  $\tilde{\Omega}(T^{\frac{2}{3}})$ . These results demonstrate that no-regret learning is possible for feature-based dynamic pricing under weak assumptions, but also reveal a disappointing fact that the seemingly richer *pricing feedback* is not significantly more useful than the *bandit-feedback* in regret reduction.

## 1 INTRODUCTION

In a dynamic pricing process, a seller presents prices for the products and adjusts these prices according to customers’ feedback (i.e., whether they decide to buy or not) to maximize the revenue. Existing works on the single-product pricing problem (Kleinberg and Leighton, 2003; Wang et al., 2021) assume that cus-

tomers make decisions only according to the comparisons between prices and their own (random) valuations, and the goal is to find out a best fixed price that maximizes the (expected) revenue. In general, the single-product pricing problem has been well studied under a variety of assumptions.

However, these methods are not applicable when there are thousands of highly differentiated products with no experience in selling them. This motivates the idea of “contextual pricing” (Cohen et al., 2020; Mao et al., 2018; Javanmard and Nazerzadeh, 2019; Liu et al., 2021), where each sale session is described by a context that also affects the valuation and pricing.

Contextual pricing. For  $t = 1, 2, \dots, T$  :

1. A context  $x_t \in \mathbb{R}^d$  is revealed that describes a sales session (product, customer and context).
2. The customer values the product as  $y_t$  using  $x_t$ .
3. The seller proposes a price  $v_t > 0$  concurrently (according to  $x_t$  and historical sales records).
4. The transaction is successful if  $v_t \leq y_t$ , i.e., the seller gets a reward  $r_t = v_t \cdot \mathbf{1}(v_t \leq y_t)$ .

Here  $T$  is the time horizon known to the seller in advance<sup>1</sup>,  $x_t$ ’s can be either stochastic (i.e., each  $x_t$  is independently and identically distributed) or adversarial (i.e., the sequence  $\{x_t\}_{t=1}^T$  are arbitrarily chosen and fixed by nature before  $t = 0$ ), and  $\mathbf{1}_t := \mathbf{1}(v_t \leq y_t)$  is an indicator that equals 1 if  $v_t \leq y_t$  and 0 otherwise. In this work, we consider two distinct problem setups that make use of the feature vector  $x_t$ .

- (a) **Linear Policy (LP)**:  $(x_t, y_t)$ ’s are selected by nature (or an oblivious adversary) arbitrarily, and the learning goal is to compete with the optimal linear prices  $v_t^* = x_t^\top \beta^*$  where  $\beta^*$  maximizes the cumulative reward in the hindsight.
- (b) **Linear Valuation (LV)**: assume *valuations* are linear + noise, i.e.,  $y_t = x_t^\top \theta^* + N_t$ , where  $\theta^* \in \mathbb{R}^d$  is a fixed vector and  $N_t$  is a market noise,

---

Proceedings of the 25<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

---

<sup>1</sup>Here we assume  $T$  known for simplicity of notations. In fact, if  $T$  is unknown, then we may apply a “doubling epoch” trick as Javanmard and Nazerzadeh (2019) and the regret bounds are the same.

drawn i.i.d. from a fixed *unknown* distribution  $\mathbb{D}$ . The learning goal is to compete with the *globally* optimal price  $v_t^* = \operatorname{argmax}_v v \cdot \Pr[v \leq y_t | x_t]$  with no restrictions on the pricing policy.

These two problem models, although quite similar at a glimpse, are intrinsically different. On the one hand, the LP problem makes no assumptions on the  $x_t \rightarrow y_t$  mapping, i.e., agnostic learning. Customers’ valuations are not necessarily linear (and can be deterministic/noisy/stochastic/adversarial), but the seller competes with the optimal policy in a constrained family. On the other hand, the LV problem makes mild modeling assumptions about the distribution of  $y_t$  given  $x_t$  while keeping the policy class unrestricted. In other words, LP is modeling our strategy while LV is modeling the nature. We adopt *regret* as a metric of algorithmic performance: For the LP problem, we compare its (expected) reward with that of the optimal fixed  $\beta^*$  in hindsight (i.e., an *ex post* regret); For the LV problem, we compare its (expected) reward with the largest expected reward condition on  $\theta^*$  and  $\mathbb{D}$  (i.e., an *ex ante* regret). We will clarify the difference between LP and LV in Appendix C.1 with more details and examples. However, we are still facing agnostic (valuation/noise) distributions in both LP and LV, and we want to know whether there exist no-regret algorithms for them.

**Summary of Results.** Our contributions are three-fold.

1. For the LP problem with adversarial  $x_t$ ’s, we present an algorithm “Linear-EXP4” that achieves  $\tilde{O}(d^{\frac{1}{3}}T^{\frac{2}{3}})$  regret.
2. For the LV problem with adversarial  $x_t$ ’s, we present an algorithm “D2-EXP4” that achieves  $\tilde{O}(T^{\frac{3}{4}} + d^{\frac{1}{2}}T^{\frac{5}{8}})$ .
3. We present an  $\tilde{O}(d^{\frac{1}{3}}T^{\frac{2}{3}})$  lower regret bound for LP problem and an  $\tilde{O}(T^{\frac{2}{3}})$  for LV problem (even with stochastic  $x_t$ ’s, known  $\theta^*$  and Lipschitz valuation distribution). The results indicate “Linear-EXP4” optimal up to logarithmic factors.

To the best of our knowledge, we are the first to study the LP problem and the version of the LV problem with no assumption on the noise. Comparing to the existing literature on this problem (Cohen et al., 2020; Javanmard and Nazerzadeh, 2019), our model makes fewer assumptions. Our results for LP is information-theoretically optimal, and our results in LV improve over the best known upper and lower bounds (from  $\tilde{O}(T^{\frac{2}{3} \vee (1-\alpha)})$  on i.i.d.  $x_t$ ’s with an indeterministic  $\alpha$  and  $\Omega(T^{\frac{3}{5}})$  in Luo et al. (2021)).

**Technical Novelty.** In this work, we make use of the *half-Lipschitz* nature in pricing problems: the proba-

bility of a price to be accepted will not decrease as the price decreases. This has been used in Kleinberg and Leighton (2003) and Cohen et al. (2020). However, they directly applied this property in discretizing the action and policy spaces, which would lead to a linear regret in our LV problem setting. In our algorithm D2-EXP4, we settle this issue by also discretizing the noise distribution space and include these discretized CDF’s as part of policy candidates. We also carefully adopt a conservative subtractor on the discretized outputs to ensure a large-enough probability of acceptance. In this way, we may get rid of any assumption on the noise distribution (even the basic Lipschitzness assumed by Luo et al. (2021)) and achieve a sub-linear regret. This discretization method, along with the conservative subtractor, can be easily transfer to any pricing problem settings with unknown i.i.d. noise. For the lower bound proof, we adapt the nested intervals and bump functions introduced by Kleinberg (2004) for continuum bandits to our pricing problem models.

## 2 RELATED WORKS

In this section, we discuss how our work relates to the existing literature on (contextual/non-contextual) pricing, bandits, and contextual search.

**Non-Contextual Dynamic Pricing.** Dynamic pricing was extensively studied under the single-product (non-contextual) setting (Kleinberg and Leighton, 2003; Besbes and Zeevi, 2009, 2012; Wang et al., 2014; Besbes and Zeevi, 2015; Chen et al., 2019; Wang et al., 2021). The crux of pricing is to learn the demand curve (i.e., the noise distribution in our LP problem) from Boolean-censored feedback. Wang et al. (2021) concludes existing results and characterizes the impact of different assumptions on the demand curve on the minimax regret. The problem of contextual dynamic pricing is more challenging mainly because we need to learn the valuation parameter  $\theta^*$  and the noise distribution jointly. Knowing one would imply a learning algorithm for another (Javanmard and Nazerzadeh, 2019; Luo et al., 2021), but learning both together makes the problem highly nontrivial.

**Contextual Dynamic Pricing.** There is a growing body of recent works focusing on the LV model of the contextual dynamic pricing problem (Cohen et al., 2020; Javanmard and Nazerzadeh, 2019; Xu and Wang, 2021; Luo et al., 2021), but most of them make strong assumptions about the noise. Table 1 lists the best existing results under these assumptions. Besides these works, Cohen et al. (2020) also achieved an  $O(d \log T)$  regret when the variance of the Sub-Gaussian noise is extremely small, i.e.,  $\tilde{O}(1/T)$ . It is worth mentioning that our “Linear-EXP4” shares the same discretization factor with “ShallowPricing” algorithm in Cohen et al.

Table 1: Related Works and Regret Bounds

Problem	Linear (Noisy) Valuation (LV)				Linear Policy (LP)
Noise Assumption	No Noise	Known, Log-concave	Parametric, Log-concave	Arbitrary bounded noise	
Upper Bound	$O(d \log \log T)$ (Leme and Schneider, 2018)	$O(d \log T)$ (Xu and Wang, 2021)	$\tilde{O}(\sqrt{T})$ (Javanmard and Nazerzadeh, 2019)	$\tilde{O}(T^{\frac{1}{2}} + d^{\frac{1}{2}} T^{\frac{1}{4}})$ (This Work)	$\tilde{O}(d^{\frac{1}{2}} T^{\frac{1}{4}})$ (This Work)
Lower Bound	$\Omega(d \log \log T)$ (Kleinberg and Leighton, 2003)	$\Omega(d \log T)$ (Javanmard and Nazerzadeh, 2019)	$\Omega(\sqrt{T})$ (Xu and Wang, 2021)	$\tilde{\Omega}(T^{\frac{1}{4}})$ (This Work)	$\tilde{\Omega}(d^{\frac{1}{2}} T^{\frac{1}{4}})$ (This Work)

(2020), but ours solves a different problem. The closest work to ours is the recent work of Luo et al. (2021) that studies the LV problem under only smoothness and concavity assumptions. They develop a UCB-style algorithm that achieves  $\tilde{O}(T^{\frac{2}{3} \vee (1-\alpha)})$  regret, assuming that a good estimator might approach  $\theta^*$  with  $O(T^{-\alpha})$  error only with the logged data. However, such an estimator was neither described nor trivial to construct with  $\alpha > 0$ . For this reason, this result does not imply a sub-linear regret bound on our problem. In comparison, our “D2-EXP4” algorithm achieves an  $\tilde{O}(T^{\frac{3}{4}})$  regret with no distributional assumptions.

**Bandits** A multi-armed bandit (MAB) is an online learning model where one can only observe the feedback of the selected action at each time. Both LP and LV can be reduced to contextual bandits (Langford and Zhang, 2007; Agarwal et al., 2014) as long as the policies and prices are finite. In this work, we make use of an “EXP-4” algorithm (Auer et al., 2002) in a new way: By carefully discretizing the parameter space and distribution functions, we enable EXP-4 agents to find out near-optimal policies among infinite continuum policy spaces. There exists another family of bandit problem: continuum-armed bandit (CAB) (Agrawal, 1995; Kleinberg, 2004; Auer et al., 2007), where the action space is continuum and the reward function is Lipschitz. In this work, we adapt the (bump functions, nested intervals) structures in Kleinberg (2004) to our lower bound proof. This adaptation is non-trivial since (1) their reward functions is not suitable for pricing problems, and (2) their feedback is not Boolean-censored.

Our results on the LP problem reveal that a reduction to contextual bandits is “tight” in regret bounds. A similar situation also occurs in Kleinberg and Leighton (2003) on non-contextual pricing. These results indicate a pricing feedback is not substantially richer than a bandit feedback in information theory, which is surprising as a pricing feedback indicates the potential feedback of a “halfspace” rather than a single point. However, does this imply we cannot get any extra information from a pricing feedback? Notice that we are matching a no-Lipschitz upper bound with a Lipschitz lower bound! In fact, a revenue curve is naturally “half Lipschitz”, which helps us get rid of this assumption. We will discuss this property in Paragraph 4.2.

**Contextual search** Contextual pricing is cohesively related to contextual search problems (Leme and Schneider, 2018; Lobel et al., 2018; Liu et al., 2021; Krishnamurthy et al., 2021) where they also learn from Boolean feedback and usually assume linear contexts. However, they are facing slightly different settings: Leme and Schneider (2018); Lobel et al. (2018) are noiseless and could achieve an optimal  $O(\log \log T)$  regret; Liu et al. (2021) allows noises directly on customers’ decisions instead of the valuations in our setting; Krishnamurthy et al. (2021) allows only small-variance valuation noises that is similar to Cohen et al. (2020).

### 3 PRELIMINARIES

**Symbols and Notations.** Now we introduce the mathematical symbols and notations involved in the following pages. The game consists of  $T$  rounds.  $x_t, \beta^*, \theta^* \in \mathbb{R}_+^d, y_t, N_t \in \mathbb{R}, v_t \in \mathbb{R}_+^2$ , where  $d \in \mathbb{Z}_+$ . At each round, we receive a payoff (reward)  $r_t = v_t \cdot \mathbb{1}_t$  where  $\mathbb{1}_t := \mathbb{1}(v_t \leq y_t)$  indicates the acceptance of  $v_t$ , i.e.,  $\mathbb{1}_t = 1$  if  $v_t \leq y_t$  and 0 otherwise. For LP problem, we denote  $F_{LP}(v|x)$  as a *demand function*, i.e. the probability of price  $v$  being accepted given feature  $x$ . Therefore,  $F_{LP}(v|x)$  is non-increasing with respect to  $v$ , for any  $x \in \mathbb{R}^d$ . For LV problem, we specifically denote  $u_t = x_t^\top \theta^*$  as the *noiseless valuation* (or *expected valuation* for zero-mean noises), and denote  $F$  as its CDF. Finally, we define  $h(v, x) = v \cdot F_{LP}(v|x)$  as an *expected revenue* function of price  $v$  given feature  $x$  in an LP problem, and  $g(v, u, F) := v \cdot (1 - F(v - u))$  as an *expected revenue* function of price  $v$  given any noiseless valuation  $u$  and noise distribution  $F$  in an LV problem.

We may use discretization methods in the following sections. Here we adopt the notation in Cohen et al. (2020) by denoting

$$\lfloor x \rfloor_\gamma := \lfloor \frac{x}{\gamma} \rfloor \cdot \gamma, \lceil x \rceil_\gamma := \lceil \frac{x}{\gamma} \rceil \cdot \gamma. \quad (1)$$

as the  $\gamma$ -lower/upper rounding of  $x$ , which discretize  $x$  as its nearest smaller/larger integer multiples of  $\gamma$ . Similarly, for  $\theta \in \mathbb{R}^d$ , we may define  $\lfloor \theta \rfloor_\gamma := [\lfloor \theta_1 \rfloor_\gamma, \lfloor \theta_2 \rfloor_\gamma, \dots, \lfloor \theta_d \rfloor_\gamma]^\top$  and  $\lceil \theta \rceil_\gamma :=$

<sup>2</sup>We do not assume  $y_t \geq 0$  since some customer would not buy anything despite the price.

$[\lceil\theta_1\rceil_\gamma, \lceil\theta_2\rceil_\gamma, \dots, \lceil\theta_d\rceil_\gamma]^\top$ . Based on this, we define a counting set  $N_{\gamma,a} := \left\{0, 1, 2, \dots, \lfloor \frac{a}{\gamma} \rfloor\right\}$ .

**Regret Definitions.** Next we define the regrets in both problems.

**Definition 1** (Regret in LP). *We define  $Reg_{LP}$  as the regret of the Linear Policy pricing problem.*

$$Reg_{LP} := \max_{\beta} \sum_{t=1}^T h(x_t^\top \beta, x_t) - h(v_t, x_t). \quad (2)$$

**Definition 2** (Regret in LV). *We define  $Reg_{LV}$  as the regret of the Linear Noisy Valuation problem.*

$$Reg_{LV} := \sum_{t=1}^T \max_v g(v, u_t, F) - g(v_t, u_t, F). \quad (3)$$

Again, we aim at competing with the best fixed  $\beta^* = \operatorname{argmax}_{\beta} \sum_{t=1}^T h(x_t^\top \beta, x_t)$  in an LP problem, and with the global best pricing policy (maximizing expected revenue at every  $t$ ) in an LV problem.

**Summary of Assumptions** We specify the problems by the following assumptions:

**Assumption 1** (bounded features and parameters). *Without losing generality, we assume that  $x_t, \beta^*, \theta^* \in \mathbb{R}_+^d, \|x_t\|_2 \leq B, \|\beta^*\|_2 \leq 1, \|\theta^*\|_2 \leq 1$ , where  $B \in \mathbb{Z}^+$  is a constant known to us in advance.*

**Assumption 2** (decreasing demand in LP). *In LP problem, assume that  $F_{LP}(v|x)$  is non-increasing for any  $v \geq 0, x \in \mathbb{R}_+^d$ .*

**Assumption 3** (bounded noise). *In LV problem, assume that  $N_t \in [-1, 1]$  that is i.i.d. sampled from a fixed unknown distribution  $\mathbb{D}$ .*

These assumptions are mild and common for algorithm design. Based on these assumptions above, we only have to consider prices in  $[0, B]$  for LP problems and  $[0, B + 1]$  for LV problems. Besides, we assume that  $T \geq d^4$  for a simplicity of comparing among different terms in regret bounds. In Section 5.2, we will introduce more assumptions to the distribution functions to demonstrate that our lower bounds hold *even if* those assumptions are made.

## 4 ALGORITHM

In this section, we propose two algorithms, Linear-EXP4 and D2-EXP4, for LP and LV problems respectively. Both of them are based on the EXP-4 algorithm (Auer et al., 2002) along with discretized policy sets. First of all, we define these policy sets:

**Definition 3** (parameter set). *For any small  $0 < \Delta < 1$ , we define a parameter set  $\Omega_{\Delta,d} \subset \mathbb{R}^d$ :*

$$\Omega_{\Delta,d} := \left\{ \|\theta\|_2 \leq 1, \theta = [n_1 \Delta, n_2 \Delta, \dots, n_d \Delta]^\top, n_1, n_2, \dots, n_d \in N_{\Delta,1} \right\}$$

**Definition 4** (CDF set). *For any small  $0 < \gamma < 1$ , we define a Cumulative Distribution Function (CDF) set  $\mathcal{F}_\gamma$ :*

$$\mathcal{F}_\gamma := \left\{ F : \mathbb{R} \rightarrow [0, 1] \text{ non decreasing, } F(v) = 0 \text{ when } v < -1, F(v) = 1 \text{ when } v > 1, \frac{F(v)}{\gamma} \in N_{\gamma,1} \text{ when } \pm \frac{v}{\gamma} \in N_{\gamma,1}, F(v) = F(\lfloor v \rfloor_\gamma) + \frac{1}{\gamma} (F(\lfloor v \rfloor_\gamma + \gamma) - F(\lfloor v \rfloor_\gamma))(v - \lfloor v \rfloor_\gamma) \text{ otherwise} \right\}$$

Definition 3 is straightforward as we use  $\Delta^d$ -grids to discretize the  $[0, 1]^d$  space. Definition 4 actually represents such a family of CDF: the random variable is defined on  $[-1, 1]$ , and its CDF equals some integer multiple of  $\gamma$  when  $v$  (or  $-v$ ) itself is an integer multiple of  $\gamma$ ; for those  $v$  in between these grids, CDF connects the two endpoints as linear. In a word, each CDF in  $\mathcal{F}_\gamma$  is a piecewise linear function with every integer-multiple- $\gamma$  points valuating some integer-multiple- $\gamma$  as well. From the definitions above, we know that  $|\Omega_{\Delta,d}| = O\left(\left(\frac{1}{\Delta}\right)^d\right)$ . Also, we have  $|\mathcal{F}_\gamma| = \binom{\frac{2}{\gamma}}{\frac{1}{\gamma}} = O\left(2^{\frac{2}{\gamma}}\right)$  according to a ‘‘balls into bins’’ model in combinatorial counting: At each point  $\frac{\pm i}{\gamma}$  (for  $i \in [\frac{2}{\gamma}]$ ) the CDF can increase by  $j \cdot \gamma$ , with  $j$  being a non-negative integer, and the summation of all increases is 1 (i.e.,  $\frac{1}{\gamma}$  of  $\gamma$  increments).

Finally we introduce the EXP-4 algorithm (Auer et al., 2002) for adversarial contextual bandits. With a finite action set  $A$  and policy set  $\Pi$ , the EXP-4 agent has a regret guarantee at  $O(\sqrt{T|A| \log |\Pi|})$  in  $T$  rounds (comparing with the optimal policy in  $\Pi$ ). The following is a simplified version of EXP-4 that illustrates its mechanism. For a more detailed introduction, please directly refer to Auer et al. (2002).

EXP-4.

**Input:** Policy set  $\Pi$ , Action set  $A$ .

Initialize each policy  $i$  with weight  $w_i$ ;

**for**  $t = 1$  **to**  $T$  **do**

Set probability  $p_j(t)$  for each action  $j$  according to weights of all policies;

Get  $a_t$  by Thompson sampling the action set  $A$  according to current probability  $\{p_j(t)\}$ ;

Receive a reward  $r_t$ ;

Construct an *Inverse Propensity Scoring (IPS)* estimator  $\hat{r}_i(t)$  for the reward of each action  $i$ .

Update weights  $w_i$ 's according to  $\hat{r}_i(t)$ .

**end for**



#### 4.1 Linear-EXP4 for LP

Here we present our ‘‘Linear-EXP4’’ algorithm for the linear policy pricing problem. It takes  $\Omega_{\Delta,d}$  as the policy set and plug it into EXP-4 algorithm, which is straightforward but significant in reducing the regret. The pseudo-code of Linear-EXP4 is summarized as Algorithm 1.

---

#### Algorithm 1 Linear-EXP4

---

**Input:** Parameter set  $\Omega_{\Delta,d}$ , Action set  $A_\gamma = \{0, \gamma, 2\gamma, \dots, \lfloor B \rfloor_\gamma\}$ , parameters  $\Delta, \gamma$ .  
 Set policy set  $\Pi_{\Delta,\gamma}^{LP} = \{\pi_\beta(x) = \lfloor x^\top \beta \rfloor_\gamma, \beta \in \Omega_{\Delta,d}\}$   
 Initialize an EXP-4 agent  $\mathcal{E}_{LP}$  with  $\Pi_{\Delta,\gamma}^{LP}, A_\gamma$ ;  
**for**  $t = 1$  **to**  $T$  **do**  
      $\mathcal{E}_{LP}$  observe  $x_t$ ;  
      $\mathcal{E}_{LP}$  choose an action (price)  $v_t$ ;  
     Receive feedback  $r_t = v_t \cdot \mathbb{1}_t$  and feed it into  $\mathcal{E}_{LP}$ ;  
**end for**

---

Here the EXP-4 agent  $\mathcal{E}_{LP}$  would approach the best policy  $\pi^*$  in  $\Pi_{\Delta,\gamma}^{LP}$  within a reasonable regret. Therefore, we have to carefully choose  $\Delta$  and  $\theta$  such that the regrets of both  $\mathcal{E}_{LP}$  and  $\pi^*$  are well bounded.

#### 4.2 Discrete-Distribution-EXP4 for LV

Here we present our ‘‘Discrete-Distribution-EXP-4’’ algorithm, or D2-EXP4 for the linear noisy valuation pricing problem. Though it originates EXP-4 as well as Linear-EXP4 above, the reduction is not as straightforward. In fact, the policy set is defined as follows:

$$\Pi_{\Delta,\gamma}^{LV} = \{\pi(x; \hat{\theta}, \hat{F}) = \max\{\lfloor x^\top \hat{\theta} \rfloor_\gamma - (B+1)\gamma + \lfloor w^*(x) \rfloor_\gamma, 0\},$$

$$\left. \begin{array}{l} \text{where } w^*(x) = \underset{w}{\operatorname{argmax}} g(u + w, x^\top \hat{\theta}, \hat{F}), \hat{\theta} \in \Omega_{\Delta,d}, \hat{F} \in \mathcal{F}_\gamma \end{array} \right\} \quad (4)$$

For each policy in  $\Pi_{\Delta,\gamma}^{LV}$ , it firstly takes a  $\hat{\theta}$  from  $\Omega_{\Delta,d}$  and a  $\hat{F}$  from  $\mathcal{F}_\gamma$ , and then generate an ‘‘optimal incremental price’’  $w^*(x)$  greedily as if they are the true parameter  $\theta^*$  and the true noise distribution  $F$ . Finally, the policy take an action (price) that is the summation of  $\gamma$ -lower roundings of  $\hat{u} = x^\top \hat{\theta}$  and  $w^*(x)$  to fit in the action set  $A_\gamma := \{0, \gamma, 2\gamma, \dots, \lfloor B+1 \rfloor_\gamma\}$ , and minus a  $(B+1)\gamma$  amount. We know that  $|\Pi_{\Delta,\gamma}^{LV}| = |\Omega_{\Delta,d}| \cdot |\mathcal{F}_\gamma| = O((\frac{1}{\Delta})^d \cdot 2^{\frac{3}{\gamma}})$ . We present the psuedo-code of D2-EXP4 as Algorithm 2.

---

#### Algorithm 2 Discrete-Distribution-EXP-4(D2-EXP4)

---

**Input:** Policy set  $\Pi_{\Delta,\gamma}^{LV}$ , Action set  $A_\gamma = \{0, \gamma, 2\gamma, \dots, \lfloor B+1 \rfloor_\gamma\}$ , parameters  $\Delta, \gamma$ .  
 Initialize an EXP-4 agent  $\mathcal{E}_{LV}$  with  $\Pi_{\Delta,\gamma}^{LV}, A_\gamma$ ;  
**for**  $t = 1$  **to**  $T$  **do**  
      $\mathcal{E}_{LV}$  observe  $x_t$ ;  
      $\mathcal{E}_{LV}$  select an action(price)  $v_t$ ;  
     Receive feedback  $r_t = v_t \cdot \mathbb{1}_t$  and feed it into  $\mathcal{E}_{LV}$ ;  
**end for**

---

D2-EXP4 is straightforward except the  $(B+1)\gamma$ -lower price than the  $\gamma$ -rounding of a greedy price. This is because we want a conservative price, and the  $(B+1)\gamma$  penalty is to compensate the ‘‘exaggerate’’  $\lfloor \theta \rfloor_\gamma$  parameter we adopt in  $\Pi_{\Delta,\gamma}^{LV}$ . We will include more details in Paragraph 4.2 below and in Section 5.1.

#### Adversarial Features and Agnostic Distributions

Notice that both algorithms are suitable for adversarial  $x_t$  series, which is a property of EXP-4. It is worth mentioning that our Linear-EXP4 makes **no** assumptions on the distribution of  $y_t$  given  $x_t$ , and that D2-EXP4 assumes **no** pre-knowledge or technical assumptions on the noise distribution (despite that noises are bounded).

#### Conservative Pricing Strategy

Both of our algorithms adopt a *conservative* strategy while pricing: In Linear-EXP4, a good-enough linear policy is the  $\gamma$ -lower rounding of parameter  $\beta^*$ ; in D2-EXP4, we even define each policy by proposing a ‘‘greedy-and-safe’’ price which is  $(B+1)\gamma$ -less than the optimal. This is because of the ‘‘half-Lipschitz’’ nature of a demand curve: decreasing the price would at least maintain the chance of being accepted. Since we do not make any Lipschitz or smoothness assumptions on the distributions, these discretizations might marginally increase the price and cause drastic change of the expected revenue. In order to avoid this, it is always better to decrease the proposed price by an acceptable small step as it guarantees the probability.

#### Computational Efficiency

Our algorithms require exponential computations w.r.t. dimension  $d$  since the EXP-4 agent requires exponential time to evaluate each policy in the policy set. An ‘‘optimization oracle’’-efficient contextual bandit algorithm in Agarwal et al. (2014) can be used in place of EXP-4 to achieve a near-optimal regret (up to logarithmic factors), but it requires the input features  $x_t$  to be drawn from an unknown fixed distribution.

## 5 REGRET ANALYSIS

In this section, we analyze our Linear-EXP4 and D2-EXP4 algorithm and prove their  $\tilde{O}(d^{\frac{1}{3}}T^{\frac{2}{3}})$  and  $O(T^{\frac{3}{4}})$

regret bounds, respectively. Also, we present a scenario where a lower bound construction with  $\tilde{\Omega}(T^{\frac{2}{3}})$  regret fits for both LP and LV problems, even under stronger assumptions including stochastic  $x_t$ 's, Lipschitz distribution functions and unimodal demand curves.

### 5.1 Upper Bounds

Here we propose the following theorem as a regret bound of Linear-EXP4. This only requires the assumption that features  $x_t$ 's and (potential) optimal parameter  $\beta^*$  is bounded by  $L_2$ -norm, without making any specifications on the feature-valuation mapping.

**Theorem 5** (Regret of Linear-EXP4). *In any LP problem, with Assumption 1, the expected regret of Linear-EXP4 does not exceed  $O(d^{\frac{1}{3}}T^{\frac{2}{3}} \log dT)$  by setting  $\Delta = T^{-\frac{1}{3}}d^{-\frac{1}{6}}$  and  $\gamma = T^{-\frac{1}{3}}d^{\frac{1}{3}}$ .*

*Proof.* We denote  $\tilde{\beta}^* = \lfloor \beta^* \rfloor_{\Delta}$  and  $\hat{\beta}^* := \operatorname{argmax}_{\beta \in \Omega_{\Delta,d}} \sum_{t=1}^T \mathbb{E}[h(\pi_{\beta}(x_t), x_t)]$ . Now we decompose the regret of LP problem as follows:

$$\begin{aligned}
 \mathbb{E}[\operatorname{Reg}_{LP}] &= \sum_{t=1}^T \mathbb{E}[h(x_t^{\top} \beta^*, x_t) - h(v_t, x_t)] \\
 &= \sum_{t=1}^T \mathbb{E}[h(x_t^{\top} \beta^*, x_t) - h(\pi_{\tilde{\beta}^*}(x_t), x_t)] \\
 &\quad + \mathbb{E}[h(\pi_{\tilde{\beta}^*}(x_t), x_t) - h(\pi_{\hat{\beta}^*}(x_t), x_t)] \\
 &\quad + \mathbb{E}[h(\pi_{\hat{\beta}^*}(x_t), x_t) - h(v_t, x_t)] \\
 &\leq \sum_{t=1}^T (x_t^{\top} \beta^* - x_t^{\top} \tilde{\beta}^*) F_{LP}(x_t^{\top} \beta^* | x_t) \\
 &\quad + \mathbb{E}[h(\pi_{\tilde{\beta}^*}(x_t), x_t) - h(\pi_{\hat{\beta}^*}(x_t), x_t)] \\
 &\quad + \mathbb{E}[h(\pi_{\hat{\beta}^*}(x_t), x_t) - h(v_t, x_t)] \\
 &\leq \sum_{t=1}^T B \cdot \Delta \sqrt{d} + 0 + \sqrt{T \cdot \frac{1}{\gamma} \cdot \log\left(\frac{1}{\Delta}\right)^d} \\
 &= O(d^{\frac{1}{3}} T^{\frac{2}{3}} \log dT).
 \end{aligned} \tag{5}$$

Here the third row is because  $\pi_{\tilde{\beta}^*}(x_t) = \lfloor x_t^{\top} \tilde{\beta}^* \rfloor_{\gamma} \leq x_t^{\top} \tilde{\beta}^* \leq x_t^{\top} \beta^*$  since  $x_t, \beta \in \mathbb{R}_+^d$  (and thus  $F_{LP}(x_t^{\top} \beta^*) \leq F_{LP}(x_t^{\top} \tilde{\beta}^*)$ ); The fourth row is because  $(x_t^{\top} \beta^* - x_t^{\top} \tilde{\beta}^*) \leq \|x_t\|_2 \cdot \|\beta^* - \tilde{\beta}^*\| \leq B \cdot \Delta \sqrt{d}$ , the optimality definition of  $\hat{\beta}^*$  and the regret bound of EXP-4 from [Auer et al. \(2002\)](#); The last row is got by plugging in the value of  $\Delta$  and  $\gamma$ . ■

The proof of Theorem 5 is straightforward based on the existing  $O(\sqrt{T|A| \log |I|})$  bound of EXP-4. We only have to bound the error of the optimal policy in  $\Pi_{\Delta, \gamma}$ . Now we present our result on D2-EXP4:

**Theorem 6** (Regret of D2-EXP4). *For any LV problem, with Assumptions 1, 2 and 3, our algorithm D2-EXP4 guarantees a regret no more than*

$O(T^{\frac{3}{4}} + T^{\frac{2}{3}} d^{\frac{1}{2}} \log dT)$  as we set  $\Delta = T^{-\frac{1}{4}} d^{-\frac{1}{2}}$  and  $\gamma = T^{-\frac{1}{4}}$ .

The proof of Theorem 6 is more sophisticated than that of Theorem 5, but they shares similar structures: we figure out one specific policy in  $\Pi_{\Delta, \gamma}^{LV}$  that is close to the optimal policy of the LV problem. The main idea of this proof is to find out a tuple of  $(\hat{\theta}, \hat{F})$  that approaches the true parameter and distribution, and to verify that the policy built on this approaching tuple is reliable only within small tractable error. The highlight is that we do not assume any Lipschitzness on the distribution, which is quite different from existing approximation methods. In fact, it is the natural property of pricing problems that enables this: for two prices  $v_1 \geq v_2$ , the probability of  $v_2$  being accepted is greater (or equal) than that of  $v_1$ , and thus  $(v_1 - v_2) \geq g(v_1, u, F) - g(v_2, u, F)$ . We may call it a *Half-Lipschitz* property since it only upper bounds the increasing rates.

Here we show a proof sketch of Theorem 6, and leave the bulk to Appendix A.

*Proof Sketch.* For any specific LV problem with linear parameter  $\theta^*$  and noise CDF  $F$ , we define  $\hat{\theta}^* := \lceil \theta^* \rceil_{\Delta}$  and  $\hat{F}$ :

$$\begin{aligned}
 \hat{F}(x) &= \lfloor F(x) \rfloor_{\gamma} \text{ when } x = i \cdot \gamma \text{ for } i \in \mathbb{Z}, \text{ and} \\
 &\text{linearly connecting } \hat{F}(i\gamma) \text{ with } \hat{F}(i+1)\gamma \tag{6} \\
 &\text{when } x \in (i\gamma, (i+1)\gamma).
 \end{aligned}$$

Our goal is to prove that  $\pi(x; \hat{\theta}^*, \hat{F})$  performs well enough. We may furthermore define a few amounts:

- (i)  $\hat{u} = x^{\top} \hat{\theta}^*$ ;
- (ii)  $w^*(u) = \operatorname{argmax}_w g(u + w, u, F)$ ;
- (iii)  $\hat{w}^*(u) = \operatorname{argmax}_w g(u + w, u, \hat{F})$ ;
- (iv)  $\hat{w}(\hat{u}) = \operatorname{argmax}_w g(\hat{u} + w, \hat{u}, \hat{F})$ .

Therefore, the price our algorithm proposed for feature  $x$  is  $\hat{v}(x) = \lfloor \hat{u} \rfloor_{\gamma} - (B+1)\gamma + \lfloor \hat{w}(\hat{u}) \rfloor_{\gamma}$ , and our goal is to prove that  $g(\hat{v}, u, F) \geq g(u + w^*(u), u, F) - C \cdot \gamma$  for some constant  $C$ . Since  $\gamma = T^{-\frac{1}{4}}$ , this would upper bounds the optimality error up to  $O(T \cdot \gamma) = O(T^{\frac{3}{4}})$ . In fact, we have the following properties:

- (i)  $\hat{\theta}^* = \lceil \theta^* \rceil_{\Delta}$  (by definition);
- (ii)  $\|\hat{\theta}^*\|_2 \leq \|\hat{\theta}^*\|_2 \leq \|\theta^*\|_2 + \Delta \sqrt{d} = \|\theta^*\|_2 + \gamma$ ;
- (iii)  $u - \gamma \leq \hat{u} - \gamma \leq \lfloor \hat{u} \rfloor_{\gamma} \leq \hat{u} \leq u + B\gamma$ ;
- (iv)  $\hat{F}(i\gamma) \leq F(i\gamma) \leq \hat{F}(i\gamma) + \gamma$ .

According to these properties, we may derive:

$$\begin{aligned}
 & g([\hat{u}]_\gamma - (B+1)\gamma + [\hat{w}(\hat{u})]_\gamma, u, F) \\
 & \geq (u + [\hat{w}(\hat{u})]_\gamma)(1 - F([\hat{w}(\hat{u})]_\gamma)) - (B+2)\gamma \\
 & \geq (u + [\hat{w}(\hat{u})]_\gamma)(1 - \hat{F}(\hat{w}(\hat{u}))) - (2B+3)\gamma \\
 & \geq g(\hat{u} + \hat{w}(\hat{u}), \hat{u}, \hat{F}) - (3B+4)\gamma \\
 & \geq g(u + \hat{w}^*(u), u, \hat{F}) - (3B+4)\gamma \\
 & \geq g(u + w^*(u), u, F) - (3B+5)\gamma.
 \end{aligned}$$

The derivation of each step is shown in Appendix A. With this policy-realizability error being bounded by  $(3B+5)\gamma = O(T^{\frac{3}{4}})$  and the original regret of the EXP-4 agent being  $O(\sqrt{TK \log N}) = \tilde{O}(T^{\frac{3}{4}} + d^{\frac{1}{2}}T^{\frac{5}{8}})$ , we may finally get a  $\tilde{O}(T^{\frac{3}{4}} + d^{\frac{1}{2}}T^{\frac{5}{8}})$  upper regret bound. ■

## 5.2 Lower Bounds

In this part, we present an  $\tilde{\Omega}(T^{\frac{2}{3}}d^{\frac{1}{3}})$  and an  $\tilde{\Omega}(T^{\frac{2}{3}})$  lower regret bounds that hold for LP and LV problems respectively. We will firstly claim a lower bound for non-contextual pricing problem, and then generalize the result to LP and LV.

**Theorem 7** (Lower bound for non-contextual pricing). *For a non-contextual pricing problem where the valuation  $y_t$ 's are generated independently and identically from a fixed unknown distribution satisfying (1) the CDF  $F(y)$  is Lipschitz and (2) the revenue curve  $g(v, F) = y \cdot (1 - F(v))$  is unimodal (i.e., non-decreasing on  $(0, v_0)$  and non-increasing on  $(v_0, +\infty)$  for some  $v_0$ ), NO algorithm can achieve  $O(T^{\frac{2}{3}-\delta})$  for any  $\delta > 0$ .*

The detailed proof of Theorem 7 is in Appendix B, and in the main pages we briefly demonstrate the constructions of the subproblem family where we achieve this lower bound.

Here we take the idea of Kleinberg (2004) where they make use of bump functions and nested intervals to ensure Lipschitz continuity and unimodality, sequentially. Since that their model is not capturing a revenue curve and that their feedback is numerical instead of Boolean, we have to adjust their design to satisfy the pricing setting. On the one hand, the probability of a price to be accepted, i.e., the rate  $\frac{\mathbb{E}[r(v)]}{v}$ , is non-increasing as the prices increases, which is not guaranteed for that of a reward function of a continuum bandit (if we treat  $v$  as an action). In this proof, we adopt a series of transformations to convert the ‘‘bump function tower’’ into a revenue curve while keeping all monotonically-increasing/decreasing intervals unchanged. On the other hand, we still use the KL-divergence to distinguish among distributions, but in a different way. As for Boolean feedback, we only need to calculate the KL-divergence of two Bernoulli random variables, which can be upper bounded by a quadratic term of their probabilistic difference.

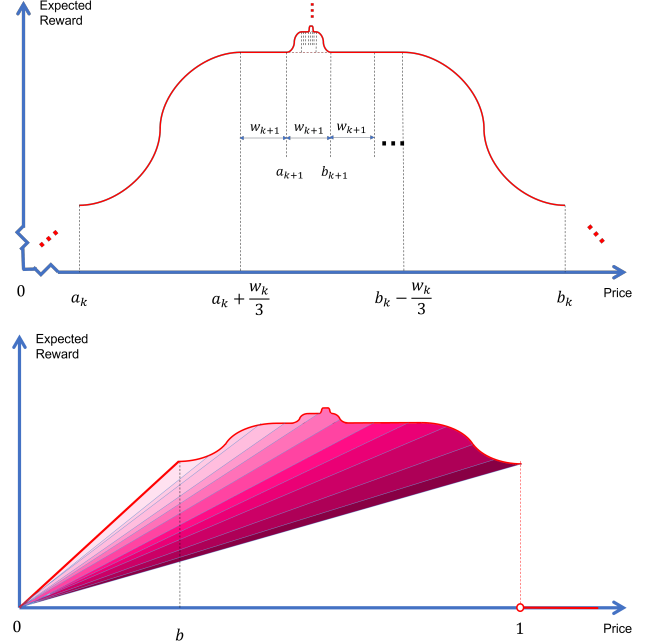


Figure 1: Structure of our lower bound function family. The upper figure shows how we use bump functions to construct a reward function  $f(v)$ . Each bump function locates at  $[a_k, b_k]$  with a length  $w_k = 3^{-k!}$ . Notice that the middle one-third of each bump is a plain divided into small intervals of length  $w_{k+1}$ , and we might randomly choose one to build up the  $(k+1)$ th bump. However, the rate  $\frac{f(v)}{v}$  that indicates the probability of  $v$  to be accepted is not necessarily non-increasing, and therefore  $f(v)$  cannot capture a revenue function for pricing. The lower figure shows an ideal revenue curve  $D(v)$  which equals  $v$  for  $v \in [0, b]$  and equals  $b + (1-b)(1 - \frac{1}{f(v)+1})$  for  $v \in (b, 1]$ . The slopes indicate that  $\frac{D(v)}{v}$  is actually non-increasing. We draw the figures with exaggeration to show the hierarchical structures better.

The constructions of bump-based revenue curves are illustrated in Figure 1. Firstly, we define a nested-interval series  $[0, 1] = [a_0, b_0] \supset [a_1, b_1] \supset \dots \supset [a_k, b_k] \supset \dots$ , where  $b_k = a_k + w_k$ ,  $w_k = 3^{-k!}$ . We let  $a_k$  be chosen from the discrete set  $\{a_{k-1} + \frac{w_{k-1}}{3} + i \cdot w_k, i = 0, 1, 2, \dots, \frac{w_{k-1}}{3w_k}\}$ . Secondly, we construct Lipschitz bump functions in each  $[a_k, b_k]$  interval, the middle one-third of which is a plain line. Thirdly, we add all these bump functions up, which forms a ‘‘tower’’ with its peak randomly generated by the series of tightening intervals  $\{[a_k, b_k]\}$ . Finally, it is transformed into a revenue curve after a series of operations.

If we treat this randomly-generated function a uniformly-distributed family of functions, then we can further prove our lower bound: On the one hand, we

prove that the feedback cannot accurately locate where the “peak of the tower” is, from the perspective of information theory. In fact, any algorithm would have a constant chance of missing the peak. On the other hand, the cost of missing a peak can be lower bounded, and thus the expected regret is as well lower bounded by their product.

With this theorem holds, we can soon get the following two corollaries:

**Corollary 8** (Lower bound of LP problem). *The lower regret bound for LP problems is  $\tilde{\Omega}(d^{\frac{1}{3}}T^{\frac{2}{3}})$ , even with stochastic features and distributional properties same as those in Theorem 7.*

*Proof.* Here we construct the following LP problem: for each time, we let  $x_t = [0, \dots, 0, 1, 0, \dots, 0]^\top$  with only the  $i_t^{\text{th}}$  element equals 1, where  $i_t$  is chosen from  $\{1, 2, \dots, d\}$  uniformly at random. As a result, the problem is split into  $d$ -subproblems with each of them a non-feature pricing problem in  $\frac{T}{d}$  rounds in expectation (since the demand function  $F_{LP}(y|x)$  can be totally different and independent for different  $x$ 's). According to Theorem 7, the lower regret bound for this problem is  $\tilde{\Omega}(d \cdot (\frac{T}{d})^{\frac{2}{3}}) = \tilde{\Omega}(d^{\frac{1}{3}}T^{\frac{2}{3}})$ . ■

**Corollary 9** (Lower bound of LV problem). *The lower regret bound for LV problems is  $\tilde{\Omega}(T^{\frac{2}{3}})$ , even with stochastic features and noise-distributional properties stated in Theorem 7.*

It is worth mentioning that the noise distribution is itself an (inversed) demand function on  $(v - u)$ , i.e., it is non-increasing as  $(v - u)$  gets larger. Based on this insight, the derivation of Corollary 9 is straightforward: any non-feature pricing problem with bounded i.i.d.  $y_t$ 's can be reduced to an LV problem up to constant coefficients. In fact, suppose  $y_t \in [a, b], 0 \leq a < b$  in a non-feature pricing problem, and then we might define an LV problem by setting  $d = 1, \theta^* = \frac{a+b}{b-a}$  and  $x_t = 1, \forall t \in \mathbb{Z}_+$  since now  $x_t^\top \theta^* + N_t \in [a, b]$ . As long as the definition of LV problem does not specify the distributional properties (besides being bounded), the distribution family in the proof of Theorem 7 can be reduced to an LV problem as well. In this way, the  $\tilde{\Omega}(T^{\frac{2}{3}})$  lower bounds are applicable to LV problems.

## 6 NUMERICAL EXPERIMENTS

In this section, we conduct numerical experiments to show the validity of Linear-EXP4. We assume  $d = 2, B = 1$  as basic parameters, and assume a Gaussian noisy valuation model i.e.,  $y_t = u_t + N_t$  where  $N_t \sim \mathcal{N}(0, \frac{1}{16})$  independently for all  $t$ . For the convenience of comparing with a fixed optimal linear




Figure 2: Regrets of Linear-EXP4 on simulated examples. The plot is on log-log scales to show the regret rate: a slope of  $\alpha$  indicates an  $O(T^\alpha)$  regret. Besides, we draw error bars with 0.95 coverage. Notice that the slope of its linear fit is 0.777, which is close to the  $\tilde{O}(T^{\frac{2}{3}})$  regret rate in theory.

policy  $\beta^*$ , we let  $u_t = J^{-1}(x_t^\top \beta^*)$  for each  $t$ , where  $J(u) = \operatorname{argmax}_v g(v, u, 1 - \Phi_{\mathcal{N}(0, \frac{1}{16})})$  is a greedy pricing function defined in Xu and Wang (2021)<sup>3</sup>. In other words, the linear price  $v_t^* = x_t^\top \beta^*$  always maximizes the expected reward for any  $t$ , and we may calculate the expected *ex ante* regret by comparing  $g(v_t, u_t, 1 - \Phi_{\mathcal{N}(0, \frac{1}{16})})$  with  $g(x_t^\top \beta^*, u_t, \Phi_{\mathcal{N}(0, \frac{1}{16})})$ . According to Hoeffding’s Inequality, the *ex post* regret that we adopt for the LP problem is only  $\tilde{O}(\sqrt{T})$  different from the *ex ante* regret. Given that the regret rate of Linear-EXP4 is  $\tilde{O}(T^{\frac{2}{3}})$ , we may ignore this difference and only show the *ex ante* regret in our experiments. Since the EXP-4 learner requires pre-knowledge on  $T$  and is not an any-time algorithm (i.e., the cumulative regret is meaningful only at  $t = T$ ), we execute Linear-EXP4 for a series of  $T = \lfloor 2^{\frac{k}{5}} \rfloor$  for  $k = 27, 28, \dots, 48$ . We repeat every experiment 20 times for each setting and then take an average. The results are shown in Figure 2

Unfortunately, we did not conduct experiments on D2-EXP4 due to the exponential time-consuming property of the EXP-4 learner along with the  $2^{T^{\frac{1}{4}}}$ -size policy set. We provide the code of D2-EXP4 in our supplementary materials.

## 7 DISCUSSION

In this section, we discuss potential extensions of this work and our conjectures on the regret of LV prob-

<sup>3</sup>They also show the existence of  $J^{-1}(v)$  by showing that  $J'(u) \in (0, 1)$ .



lems.

**From Linear to Non-Linear** Both LP and LV problems are based on a linear principle of feature-price/valuation relationships, which is not reasonable in many real-world situations (for example, the price of a diamond). Based on our specifications on LP and LV problems, we may similarly define two corresponding problems: (1) We make no assumptions on the  $x_t \rightarrow y_t$  mapping, but compare with the optimal policy in a parametric non-linear model space. (2) We directly assume that the  $x_t \rightarrow y_t$  is a parametric non-linear function adding some unknown (and non-parametric) noise, and compare with the optimal price. We may slightly modify our Linear-EXP4 and D2-EXP4 to deal with these two problems by just replacing the linear discretized policy set with another non-linear one. However, we should be careful about any discretization involved: the  $\gamma$ -roundings of non-linear policy parameters do not necessarily lead to a slightly lower price (maybe either higher or much lower). Like what we designed in D2-EXP4, we still have to ensure the parametric optimal policy itself performs within a  $[-O(\gamma), 0]$  range from the global optimal policy.

**The Minimax Regret(s) of LV** Existing works on solving LV have achieved various regret bounds with different assumptions. This is quite different from the linear regression problem where noise distributions do not significantly affect the result. To the best of our knowledge, we are the first to get rid of all assumptions (despite bounded-noise assumption<sup>4</sup>). However, we did not close the regret gap in this setting. This problem is similar to a non-feature pricing problem as we adopt the same lower bound proof in this work, but the situations are entirely different: In non-feature pricing, we aim at a fixed optimal price, and we only have to know the valuation distribution around the optimal price. However, in an LV problem, we have to approach the exact linear valuation adding an optimal *increment* for each feature, and the optimal increments are **not** fixed for different valuations. As a result, we have to know the whole noise distribution. This drastically increases the hardness of LV, and we conjecture LV with a  $\Theta(T^\alpha)$  regret where  $\alpha > \frac{2}{3}$ .

**Dependence on Noise Scale  $R$**  In this work we assume the noise  $N_t \in [-1, 1]$ . Based on this assumption, we construct a discrete noise CDF family  $\mathcal{F}_\gamma$  whose size is  $\left(\frac{3}{\gamma}\right)$ . When it changes to  $N_t \in [-R, R]$  for larger  $R$ , the number of discrete CDF is  $\left(\frac{2R+1}{\gamma}\right) \leq \left(\frac{2R+1}{\gamma}\right)^{\frac{1}{\gamma}}$ . Also, this would increase the upper bound of prices

from  $(B+1)$  to  $(B+R)$ , which would increase the number of actions by  $\frac{R}{\gamma}$ . Recall that the regret of EXP-4 is  $O(\sqrt{KT \log N})$  where  $K$  is the number of actions and  $N$  is the number of policies (i.e.,  $\#$  discrete  $\theta$  times  $\#$  discrete CDF). Therefore, the dependence on  $R$  is  $O(\sqrt{R \log R})$ .

## 8 CONCLUSION

In this work, we have studied two agnostic feature-based dynamic pricing problems: a linear pricing policy (LP) problem with no assumptions on feature-valuation mappings, and a linear noisy valuation (LV) problem with agnostic noise distributions. For the LP problem, we have presented a *Linear-EXP4* algorithm whose  $\tilde{O}(T^{\frac{2}{3}} d^{\frac{1}{3}})$  regret matches the  $\tilde{\Omega}(T^{\frac{2}{3}} d^{\frac{1}{3}})$  lower bound up to logarithmic factors. For the LV problem, we have proposed an  $\tilde{O}(T^{\frac{3}{4}})$ -regret algorithm *D2-EXP4* along with an  $\tilde{\Omega}(T^{\frac{2}{3}})$  lower bound proof even with stochastic, Lipschitz and unimodal assumptions, and both of them substantially improve existing results from  $O(T^{\frac{2}{3} \cup (1-\alpha)})$  (with smoothness assumptions and indeterministic  $\alpha$ ) and  $\Omega(T^{\frac{2}{3}})$  respectively. Both Linear-EXP4 and D2-EXP4 allow adversarial features. Besides, we have discussed the prospective generalization of this work and development of future research in feature-based dynamic pricing.

## Acknowledgements

The work is partially supported by the Adobe Data Science Award and a start-up grant from the UCSB Department of Computer Science.

## References

- Agarwal, A., Hsu, D., Kale, S., Langford, J., Li, L., and Schapire, R. (2014). Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning (ICML-14)*, pages 1638–1646.
- Agrawal, R. (1995). The continuum-armed bandit problem. *SIAM journal on control and optimization*, 33(6):1926–1951.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77.
- Auer, P., Ortner, R., and Szepesvári, C. (2007). Improved rates for the stochastic continuum-armed bandit problem. In *International Conference on Computational Learning Theory*, pages 454–468. Springer.
- Bartók, G., Foster, D. P., Pál, D., Rakhlin, A., and Szepesvári, C. (2014). Partial monitoring—classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 39(4):967–997.

<sup>4</sup>If the noise is neither bounded nor parametrized, then any finite-time algorithm will suffer a linear regret when the noise is very large and prices are always being accepted.

- Besbes, O. and Zeevi, A. (2009). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420.
- Besbes, O. and Zeevi, A. (2012). Blind network revenue management. *Operations research*, 60(6):1537–1550.
- Besbes, O. and Zeevi, A. (2015). On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Science*, 61(4):723–739.
- Chen, Q., Jasin, S., and Duenyas, I. (2019). Nonparametric self-adjusting control for joint learning and optimization of multiproduct pricing with finite resource capacity. *Mathematics of Operations Research*, 44(2):601–631.
- Cohen, M. C., Lobel, I., and Paes Leme, R. (2020). Feature-based dynamic pricing. *Management Science*, 66(11):4921–4943.
- Javanmard, A. and Nazerzadeh, H. (2019). Dynamic pricing in high-dimensions. *The Journal of Machine Learning Research*, 20(1):315–363.
- Kleinberg, R. (2004). Nearly tight bounds for the continuum-armed bandit problem. *Advances in Neural Information Processing Systems*, 17:697–704.
- Kleinberg, R. and Leighton, T. (2003). 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.
- Krishnamurthy, A., Lykouris, T., Podimata, C., and Schapire, R. (2021). 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.
- Langford, J. and Zhang, T. (2007). The epoch-greedy algorithm for contextual multi-armed bandits. In *Advances in Neural Information Processing Systems (NIPS-07)*, pages 817–824.
- Leme, R. P. and Schneider, J. (2018). Contextual search via intrinsic volumes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS-18)*, pages 268–282. IEEE.
- Liu, A., Leme, R. P., and Schneider, J. (2021). Optimal contextual pricing and extensions. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA-21)*, pages 1059–1078. SIAM.
- Lobel, I., Leme, R. P., and Vladu, A. (2018). Multi-dimensional binary search for contextual decision-making. *Operations Research*, 66(5):1346–1361.
- Luo, Y., Sun, W. W., et al. (2021). Distribution-free contextual dynamic pricing. *arXiv preprint arXiv:2109.07340*.
- Mao, J., Leme, R. P., and Schneider, J. (2018). Contextual pricing for lipschitz buyers. In *NeurIPS*, pages 5648–5656.
- Wang, Y., Chen, B., and Simchi-Levi, D. (2021). Multimodal dynamic pricing. *Management Science*.
- Wang, Z., Deng, S., and Ye, Y. (2014). Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Operations Research*, 62(2):318–331.
- Xu, J. and Wang, Y.-X. (2021). Logarithmic regret in feature-based dynamic pricing. *arXiv preprint arXiv:2102.10221*.

## Appendix

### A Proof of Regret of D2-EXP4: Theorem 6

*Proof.* For any specific LV problem that is defined with linear parameter  $\theta^*$  and noise CDF  $F$ , we define another parameter  $\hat{\theta}^* := \lceil \theta^* \rceil_\Delta$  and another CDF functions  $\hat{F}$ :

$$\begin{aligned} \hat{F}(x) &= \lfloor F(x) \rfloor_\gamma \text{ when } x = i \cdot \gamma \text{ for } i \in \mathbb{Z}, \text{ and linearly connecting} \\ &\hat{F}(i\gamma) \text{ with } \hat{F}(i+1)\gamma \text{ when } x \in (i\gamma, (i+1)\gamma). \end{aligned}$$

Notice that  $\hat{F} \in \mathcal{F}_\gamma, \hat{\theta}^* \in \Omega_{\Delta, d}$ , and our goal is to prove that  $\pi(x; \hat{\theta}^*, \hat{F})$  is good enough to match the regret. With these two definitions, we might furthermore define a few amounts:  $\hat{u} = x^\top \hat{\theta}^*, w^*(u) = \operatorname{argmax}_w g(u + w, u, F), \hat{w}^*(u) = \operatorname{argmax}_w g(u + w, u, \hat{F}), \hat{w}(\hat{u}) = \operatorname{argmax}_w g(\hat{u} + w, \hat{u}, \hat{F})$ . Therefore, the price our algorithm proposed for feature  $x$  is  $\hat{v}(x) = \lfloor \hat{u} \rfloor_\gamma - (B+1)\gamma + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma$ , and our goal is to prove that  $g(\hat{v}, u, F) \geq g(u + w^*(u), u, F) - C \cdot \gamma$  for some constant  $C$ . Since  $\hat{\theta}^* := \lceil \theta^* \rceil_\Delta$ , we have  $\|\theta^*\|_2 \leq \|\hat{\theta}^*\|_2 \leq \|\theta^*\|_2 + \Delta\sqrt{d} = \|\theta^*\|_2 + \gamma$  and thus  $u - \gamma \leq \hat{u} - \gamma \leq \lfloor \hat{u} \rfloor_\gamma \leq \hat{u} \leq u + B\gamma$ . Based on this, we may get rid of  $\lfloor \hat{u} \rfloor_\gamma$  as follows:

$$\begin{aligned} &g(\lfloor \hat{u} \rfloor_\gamma - (B+1)\gamma + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma, u, F) \\ &= (\lfloor \hat{u} \rfloor_\gamma - (B+1)\gamma + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma) \cdot (1 - F(\lfloor \hat{u} \rfloor_\gamma - (B+1)\gamma + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma - u)) \\ &\geq (\lfloor \hat{u} \rfloor_\gamma + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma) \cdot (1 - F(\lfloor \hat{u} \rfloor_\gamma - (u + (B+1)\gamma) + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma)) - (B+1)\gamma \\ &\geq (u + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma)(1 - F(\lfloor \hat{w}(\hat{u}) \rfloor_\gamma)) - (B+2)\gamma. \end{aligned}$$

Now we target at  $\lfloor \hat{w}(\hat{u}) \rfloor_\gamma$  that occurs in both of the price term and the probability term, and we will get rid of it by two steps. Since  $\hat{F}(i\gamma) \leq F(i\gamma) \leq \hat{F}(i\gamma) + \gamma$ , we have the first step like:

$$\begin{aligned} &(u + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma)(1 - F(\lfloor \hat{w}(\hat{u}) \rfloor_\gamma)) \\ &\geq (u + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma)(1 - \hat{F}(\lfloor \hat{w}(\hat{u}) \rfloor_\gamma) - \gamma) \\ &\geq (u + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma)(1 - \hat{F}(\lfloor \hat{w}(\hat{u}) \rfloor_\gamma)) - (B+1)\gamma \\ &\geq (u + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma)(1 - \hat{F}(\hat{w}(\hat{u}))) - (B+1)\gamma. \end{aligned}$$

Here the second inequality comes from the  $(B+1)$  natural bound of any price. Again, we apply  $u \geq \hat{u} - B\gamma$  and get the second step:

$$\begin{aligned} &(u + \lfloor \hat{w}(\hat{u}) \rfloor_\gamma)(1 - \hat{F}(\hat{w}(\hat{u}))) \\ &\geq (\hat{u} - B\gamma + \hat{w}(\hat{u}) - \gamma)(1 - \hat{F}(\hat{w}(\hat{u}))) \\ &\geq (\hat{u} + \hat{w}(\hat{u}))(1 - \hat{F}(\hat{w}(\hat{u}))) - (B+1)\gamma \\ &= g(\hat{u} + \hat{w}(\hat{u}), \hat{u}, \hat{F}) - (B+1)\gamma. \end{aligned}$$

Now, there are only  $\hat{\cdot}$ 's instead of  $\gamma$ -roundings, and we will get rid of those  $\hat{\cdot}$ 's within some  $C \cdot \gamma$  errors. According to the definition of  $\hat{w}(\hat{u})$  that it optimizes  $g(\hat{u} + w, \hat{u}, \hat{F})$ , we further have:

$$\begin{aligned}
 & g(\hat{u} + \hat{w}(\hat{u}), \hat{u}, \hat{F}) \\
 & \geq g(\hat{u} + \hat{w}^*(u), \hat{u}, \hat{F}) \\
 & = (\hat{u} + \hat{w}^*(u))(1 - \hat{F}(\hat{w}^*(u))) \\
 & \geq (u + \hat{w}^*(u))(1 - \hat{F}(\hat{w}^*(u))) \\
 & = g(u + \hat{w}^*(u), u, \hat{F})
 \end{aligned}$$

Finally, according to the definition of  $\hat{w}^*(u)$  that it optimizes  $g(u + w, u, \hat{F})$ , we have:

$$\begin{aligned}
 & g(u + \hat{w}^*(u), u, \hat{F}) \\
 & \geq g(u + \lfloor w^*(u) \rfloor_\gamma, u, \hat{F}) \\
 & = (u + \lfloor w^*(u) \rfloor_\gamma)(1 - \hat{F}(\lfloor w^*(u) \rfloor_\gamma)) \\
 & \geq (u + \lfloor w^*(u) \rfloor_\gamma)(1 - F(\lfloor w^*(u) \rfloor_\gamma)) \\
 & \geq (u + w^*(u) - \gamma)(1 - F(w^*(u))) \\
 & \geq g(u + w^*(u), u, F) - \gamma.
 \end{aligned}$$

Here the fourth line is again due to  $\hat{F}(i\gamma) \leq F(i\gamma) \leq \hat{F}(i\gamma) + \gamma$  and the non-decreasing property of  $F$ . We make a tricky use of  $\lfloor \cdot \rfloor_\gamma$  as a ‘‘ladder’’ helping us climb between  $F$  and  $\hat{F}$ , and the ladders only emerge on those  $i\gamma$  places as  $i \in \mathbb{Z}$ . Therefore, we have  $g(\hat{v}, u, F) \geq g(u + w^*(u), u, F) - (3B + 5) \cdot \gamma$ . Since  $\gamma = T^{-\frac{1}{4}}$ , this would upper bound the optimality error up to  $O(T \cdot \gamma) = O(T^{\frac{3}{4}})$ . Also, the EXP-4 agent would cause a regret of  $O(\sqrt{T|A| \log \Pi_{\Delta, \gamma}^{LV}}) = O(\sqrt{\frac{T}{\gamma^2} + \frac{Td \log dT}{\gamma}}) = O(T^{\frac{3}{4}} + T^{\frac{5}{8}}d^{\frac{1}{2}} \log dT)$ . This completes the proof.  $\blacksquare$

## B Proof of Lower Bound: Theorem 7

Before the proof begins, we make some necessary definitions. First of all, define a *bump function* as following:

**Definition 10** (Bump function). *For  $v \in \mathbb{R}+$ , we define*

$$B(v) = \begin{cases} 0 & v \in (-\infty, 0] \cup [1, +\infty) \\ \exp\{\frac{1}{(3v-1)^2-1}\} & v \in (0, 1/3) \\ 1 & v \in [1/3, 2/3] \\ \exp\{\frac{1}{(3v-2)^2-1}\} & v \in (2/3, 1) \\ 0 & v \in [1, +\infty) \end{cases}$$

as a basic bump function. Then we define a rescaled bump function:

$$B_{[a,b]}(v) = B\left(\frac{v-a}{b-a}\right).$$

Here we present a lemma on the Lipschitzness of  $B(v)$ :

**Lemma 11** (Lipschitz continuity of  $B(v)$ ).  *$B(v)$  is 6-Lipschitz, i.e.,  $|B'(v)| \leq 6$ . Also,  $|B'_{[a,b]}(v)| = |\frac{1}{b-a} B'(\frac{v-a}{b-a})| \leq \frac{6}{b-a}$ .*

*Proof.* According to Definition 10, we have:

$$B'(v) = \begin{cases} 0 & v \in (-\infty, 0] \\ -\frac{1}{((3v-1)^2-1)^2} \cdot 6(3v-1) \exp\{\frac{1}{(3v-1)^2-1}\} & v \in (0, 1/3) \\ 0 & v \in [1/3, 2/3] \\ -\frac{1}{((3v-2)^2-1)^2} \cdot 6(3v-2) \exp\{\frac{1}{(3v-2)^2-1}\} & v \in (2/3, 1) \\ 0 & v \in [1, +\infty) \end{cases}$$

Now we propose a lemma:



**Lemma 12.** For  $t > 1$ , we have  $\frac{t^2}{e^t} \leq 1$ .

In fact, for both  $1 < t \leq \sqrt{e}$  and  $t \geq 2$ , the inequality is trivial. For  $t \in (\sqrt{e}, 2)$ , we have  $\ln(e^t) > \ln(e^{\sqrt{e}}) = \sqrt{e} \cdot 1 > 1.6 > 2 \times 0.7 > 2 \ln 2 > 2 \ln t = \ln(t^2)$ .

Now we denote  $t_1 = -\frac{1}{(3v-1)^2-1}$ ,  $t_2 = -\frac{1}{(3v-2)^2-1}$ , and we know that  $t_1 > 1$  for  $v \in (0, \frac{1}{3})$  and  $t_2 > 1$  for  $v \in (\frac{2}{3}, 1)$

$$B'(v) = \begin{cases} 0 & v \in (-\infty, 0] \\ -t_1^2 \cdot 6(3v-1) \exp\{-t_1\} & v \in (0, 1/3) \\ 0 & v \in [1/3, 2/3] \\ -t_2^2 \cdot 6(3v-2) \exp\{-t_2\} & v \in (2/3, 1) \\ 0 & v \in [1, +\infty) \end{cases}$$

Given the lemma above, we can immediately see that  $-6 \leq B'(v) \leq 6$ . This ends the proof of Lemma 11. ■

Secondly, we define a series of intervals  $[0, 1] = [a_0, b_0] \supset [a_1, b_1] \supset \dots \supset [a_k, b_k] \supset \dots$ , where  $b_k = a_k + w_k$ ,  $w_k = 3^{-k!}$ . Notice that  $w_k$  shrinks even faster than exponential series. Now we describe how to choose  $[a_k, b_k]$  from  $[a_{k-1}, b_{k-1}]$ : We divide the range  $[a_{k-1} + \frac{w_{k-1}}{3}, b_{k-1} + \frac{w_{k-1}}{3}]$  into  $Q_k = \frac{w_{k-1}}{3w_k}$  sub-intervals of the same length  $w_k$ , and then we pick one of these sub-intervals uniformly at random and denote it as  $[a_k, b_k]$ . It is trivial to see that  $[a_1, b_1] = [\frac{1}{3}, \frac{2}{3}]$ ,  $[a_2, b_2] = [\frac{4}{9}, \frac{5}{9}]$ .

Thirdly, we define a function:

$$f(v) := C_f \cdot \sum_{k=0}^{\infty} w_k \cdot B_{[a_k, b_k]}(v), \quad (7)$$

where  $C_f > 0$  is a constant which we will determine later. There are a few properties of  $f(v)$  shown in the following lemma:

**Lemma 13.** Define  $f(v)$  as Equation 7, and we have:

1. There exists a unique  $v^* \in [0, 1]$  such that  $f(v^*) = \max_{v \in [0, 1]} f(v)$ . In specific,  $v^* = \bigcap_{k=1}^{\infty} [a_k, b_k]$ .
2.  $f(v)$  is unimodal.
3. For any  $v \in [0, 1]$ , there exists at most one  $k$ , such that  $B'_{[a_k, b_k]}(v) \neq 0$ .
4.  $f(v) \leq \frac{3}{2} C_f$ .

*Proof.* To prove 1, we first see that  $v^* \in [a_k, b_k]$ ,  $k = 1, 2, \dots$ . Notice that  $\lim_{k \rightarrow \infty} a_k$  exists (since  $\{a_k\}_{k=0}^{\infty}$  is increasing and upper bounded) and that  $\lim_{k \rightarrow \infty} (b_k - a_k) = \lim_{k \rightarrow \infty} 3^{-k!} = 0$ . Therefore,  $\bigcap_{k=1}^{\infty} [a_k, b_k]$  is a unique real number within  $[\frac{1}{3}, \frac{2}{3}]$ .

To prove 2, notice that every  $B_{[a_k, b_k]}(v)$  is non-decreasing in  $[0, v^*]$  and non-increasing in  $[v^*, 1]$ .

To prove 3, consider the case when  $B'_{[a_k, b_k]}(v) \neq 0$ , and we know that: (1)  $v \in [a_k, b_k] \subset [a_{k-1} + \frac{w_{k-1}}{3}, b_{k-1} - \frac{w_{k-1}}{3}] \subset [a_{k-2} + \frac{w_{k-2}}{3}, b_{k-2} - \frac{w_{k-2}}{3}] \subset \dots \subset [a_0 + \frac{w_0}{3}, b_0 - \frac{w_0}{3}] = [\frac{1}{3}, \frac{2}{3}]$ . Since  $B'_{[a_j, b_j]}(v) = 0$ ,  $v \in [a_j + \frac{w_j}{3}, b_j - \frac{w_j}{3}]$ , we know that  $B'_{[a_j, b_j]}(v) = 0$ ,  $j = 0, 1, \dots, k-1$ . (2)  $v \notin [a_{k+1}, b_{k+1}] \supset [a_{k+2}, b_{k+2}] \supset \dots$ , and we know that  $B'_{[a_i, b_i]}(v) = 0$ ,  $i = k+1, k+2, \dots$

To prove 4, just notice that  $B(v) \leq 1$  and thus  $f(v) \leq C_f \cdot \sum_{k=0}^{\infty} 3^{-k!} \leq C_f \cdot \sum_{k=0}^{\infty} 3^{-k!} = \frac{3}{2} C_f$ . ■

According to Lemma 13 Property 3, we have:

$$\begin{aligned} |f'(v)| &\leq C_f \max_{v \in [a_k, b_k], k=1, 2, \dots} |w_k \cdot B'_{[a_k, b_k]}(v)| \\ &= C_f \max_{v \in [a_k, b_k], k=1, 2, \dots} |w_k \cdot B'(\frac{v - a_k}{b_k - a_k}) \cdot \frac{1}{w_k}| \\ &\leq C_f \max_y |B'(y)| \\ &\leq 6C_f \end{aligned}$$

This holds for any  $v \in [0, v^*) \cup (v^*, 1]$ . Now we define another function  $G(v)$ <sup>5</sup>:

$$G(v) = 1 - \frac{1}{f(v) + 1}, v \in [0, 1]. \quad (8)$$

According to Lemma 13 that reveals the properties of  $f(v)$ , we have a similar lemma on  $G(v)$ :

**Lemma 14.** *Define  $G(v)$  as Equation 8, and we have the following properties:*

1.  $G(0) = 0, G(1) = 0, 0 < G(v) < 1$  for  $v \in (0, 1)$ .
2.  $G(v)$  is unimodal in  $[0, 1]$ .
3.  $G'(v) = \frac{f'(v)}{(f(v)+1)^2} \Rightarrow |G'(v)| \leq 6C_f$ .

The proof of Lemma 14 is trivial.

Notice that  $G(v)$  is not necessarily a revenue curve, since  $\frac{G(v)}{v}$  is not necessarily decreasing (and thus not a “survival function”). However, we can construct a revenue curve  $D(v) : [0, 1] \rightarrow [0, 1]$  via an affine transformation:

$$D(v) = \begin{cases} v & v \in [0, b] \\ b + (1-b)G(\frac{v-b}{1-b}) & v \in (b, 1]. \end{cases} \quad (9)$$

Here  $b = \frac{6C_f+1}{2} \in (0, 1)$ , and therefore  $C_f < \frac{1}{6}$ . An illustration of the transformation from  $f(v)$  (the upper figure) to  $D(v)$  (the lower figure) is shown in Figure 1. The monotonicity in each interval is not changed, while the rate of  $\frac{\mathbb{E}[r(v)]}{v}$  is non-increasing after these transformations. As is mentioned above, the homothetic transformation with center  $(1, 1)$  ensures a non-increasing property of  $\frac{D(v)}{v}$ . Here we denote  $d(v) := \frac{D(v)}{v}$ . To show that  $D(v)$  is a revenue curve, we expand the definition of  $d(v)$  to  $\mathbb{R}$  as follows:

$$d(v) = \begin{cases} 1 & v \in (-\infty, 0] \\ \frac{D(v)}{v} & v \in (0, 1) \\ 0 & v \in [1, +\infty). \end{cases} \quad (10)$$

Now we claim that there exists a random variable  $Y > 0$  such that  $d(v) = \mathbb{P}[Y \geq v]$ . To show this, it is sufficient to prove the following lemma:

**Lemma 15.** *For  $d(v)$  defined in Equation 10, we have the following properties:*

1.  $d(v)$  is non-increasing on  $\mathbb{R}$ .
2.  $d(v)$  is continuous at  $v = 0$ , i.e.  $d(0) = \lim_{v \rightarrow 0^+} d(v) = 1$ .
3.  $d(v) \geq 0, v \in [0, 1]$ .

*Proof.* According to Equation 9 and 10, we have:

$$d(v) = \begin{cases} 1 & v \in (-\infty, b] \\ \frac{b}{v} + \frac{1-b}{v} \cdot G(\frac{v-b}{1-b}) & v \in (b, 1) \\ 0 & v \in [1, +\infty). \end{cases}$$

Therefore, we take the derivatives of  $d(v)$  and get:

$$\frac{\partial d(v)}{v} = \begin{cases} 0 & v \in (-\infty, 0) \cup (0, b) \\ -\frac{b-v \cdot G'(\frac{v-b}{1-b}) + (1-b)G'(\frac{v-b}{1-b})}{v^2} & v \in (b, 1) \\ 0 & v \in (1, +\infty). \end{cases}$$

<sup>5</sup>Here G stands for “gain”, which is different from the revenue curve to be introduced later.

Notice that

$$\begin{aligned}
 b - v \cdot G'\left(\frac{v-b}{1-b}\right) &\geq b - |v| \cdot \left|G'\left(\frac{v-b}{1-b}\right)\right| \\
 &\geq b - 1 \cdot 6C_f \\
 &= \frac{6C_f + 1}{2} - 6C_f \\
 &= \frac{1 - 6C_f}{2} \\
 &> 0.
 \end{aligned}$$

The last inequality comes from the fact that  $C_f < \frac{1}{6}$ . Therefore,  $d'(v)$  is non-positive in  $(-\infty, 0) \cup (0, b) \cup (b, 1) \cup (1, +\infty)$ . Also,  $d(v)$  is continuous at  $v = 0, v = b$  and  $\lim_{v \rightarrow 1^-} d(v) = b > 0 = \lim_{v \rightarrow 1^+} d(v)$ , we know that  $d(v)$  is always non-increasing on  $\mathbb{R}$ .  $\blacksquare$

Notice that there is a bijection between each  $\{[a_k, b_k]\}_{k=0}^\infty$  series and each  $f(v)$ , and correspondingly each  $G(v)$ ,  $D(v)$  and  $d(v)$ . Still, a bijection lies between each  $d(v)$  and the distribution  $\mathbb{P}[Y \geq v]$  of the customers' valuation. Therefore, we will take  $d(v)$  to represent this distribution.

With all preparations done above, we are now able to prove Theorem 7. Specifically, we will proof the theorem on an infinite series of  $n_1, n_2, \dots$ , where  $n_k = \lceil \frac{1}{k} (\frac{w_{k-1}}{w_k^3}) \rceil$ . Consider the possible  $Q_k = \frac{w_{k-1}}{3w_k}$  choices of  $[a_k, b_k]$ , and denote these intervals as  $I_j, j = 1, 2, \dots, Q_k$ . If  $[a_k, b_k] = I_j$ , then we denote the corresponding  $f(v), G(v), D(v), d(v)$  functions as  $f_j(v), G_j(v), D_j(v)$  and  $d_j(v)$  sequentially. Meanwhile, if we *do not* make any choice of  $[a_k, b_k]$ , and then we just have a finite series of intervals  $[0, 1] = [a_0, b_0] \supset [a_1, b_1] \supset [a_2, b_2] \supset \dots \supset [a_{k-1}, b_{k-1}]$ , and then we can define a  $f_0(v) = C_f \cdot \sum_{j=0}^{k-1} w_j \cdot B_{[a_j, b_j]}(v)$ , and can also define corresponding  $G_0(v), D_0(v), d_0(v)$  based on  $f_0(v)$ .

Now, consider the pricing feedbacks in total  $n$  rounds (where we denote  $n_k$  as  $n$  for simplicity). Define a *feedback vector*  $\mathbf{r}_n \in \{0, 1\}^n$ , denoting the outcome of a deterministic policy interacting with the revenue curve. We claim that for  $t = 1, 2, \dots, n$ , a vector  $\mathbf{r}_t$  is sufficient for any deterministic policy to generate a price  $v_{t+1}$ , because  $\mathbf{r}_i$  is a prefix of  $\mathbf{r}_j$  when  $i \leq j$ . For any policy  $\pi$ , denote the probability of  $\mathbf{r}_n$ 's occurrence as  $\mathbb{P}_j(\mathbf{r}_n)$  under the distribution  $d_j$ , or  $\mathbb{P}_0(\mathbf{r}_n)$  under the distribution  $d_0$ . Denote the series of prices that  $\pi$  has generated as  $\{v_t\}, t = 1, 2, \dots, n$ , and we may assume  $v_t \geq b$  without losing generality (as  $0 \leq v < b$  is always suboptimal). Then, for any function  $h : \{0, 1\}^n \rightarrow [0, M]$ , we have:

$$\begin{aligned}
 &\mathbb{E}_{\mathbb{P}_j}[h(\mathbf{r}_n)] - \mathbb{E}_{\mathbb{P}_0}[h(\mathbf{r}_n)] \\
 &= \sum_{\mathbf{r}_n} h(\mathbf{r}_n) \cdot (\mathbb{P}_j[\mathbf{r}_n] - \mathbb{P}_0[\mathbf{r}_n]) \\
 &\leq \sum_{\mathbf{r}_n: \mathbb{P}_j[\mathbf{r}_n] \geq \mathbb{P}_0[\mathbf{r}_n]} h(\mathbf{r}_n) (\mathbb{P}_j[\mathbf{r}_n] - \mathbb{P}_0[\mathbf{r}_n]) \\
 &\leq M \cdot \sum_{\mathbf{r}_n: \mathbb{P}_j[\mathbf{r}_n] \geq \mathbb{P}_0[\mathbf{r}_n]} h(\mathbf{r}_n) (\mathbb{P}_j[\mathbf{r}_n] - \mathbb{P}_0[\mathbf{r}_n]) \\
 &= \frac{M}{2} \|\mathbb{P}_j - \mathbb{P}_0\|_1 \\
 &\leq \frac{M}{2} \sqrt{2 \ln 2 \cdot KL(\mathbb{P}_0 \| \mathbb{P}_j)}.
 \end{aligned}$$

The last line comes from Lemma 11.6.1 in Cover & Thomas, Elements of Information Theory, where  $KL$  stands for the KL-divergence. Since

$$\begin{aligned}
 KL(\mathbb{P}_0(\mathbf{r}_n) \| \mathbb{P}_j(\mathbf{r}_n)) &= \sum_{t=1}^n KL(\mathbb{P}_0[r_t | \mathbf{r}_{t-1}] \| \mathbb{P}_j[r_t | \mathbf{r}_{t-1}]) \\
 &= \sum_{t=1}^t \mathbb{P}_0\left(\frac{v_t - b}{1 - b} \notin I_j\right) \cdot 0 + \mathbb{P}_0\left(\frac{v_t - b}{1 - b} \in I_j\right) \cdot KL\left(\frac{D_0(v_t)}{v_t} \parallel \frac{D_j(v_t)}{v_t}\right).
 \end{aligned}$$

The first equality comes from the chain rule of decomposing a KL-divergence. The second equality is because  $r_t$  is a Bernoulli random variable that satisfies  $Ber(\frac{D_0(v_t)}{v_t})$  under  $\mathbb{P}_0$ , or  $Ber(\frac{D_j(v_t)}{v_t})$  under  $\mathbb{P}_j$ . Denote  $\mu_t := \frac{v_t - b}{1 - b}$  for simplicity. Notice that if  $v_t \in I_j$ , then we have:

$$\begin{aligned}
 \frac{D_j(v_t)}{v_t} - \frac{D_0(v_t)}{v_t} &= \frac{(1-b) \cdot G_j(\mu_t)}{v_t} - \frac{(1-b) \cdot G_0(\mu_t)}{v_t} \\
 &= \frac{1-b}{v_t} (G_j(\mu_t) - G_0(\mu_t)) \\
 &= \frac{1-b}{v_t} \left( \frac{1}{f_0(\mu_t) + 1} - \frac{1}{f_j(\mu_t) + 1} \right) \\
 &= \frac{1-b}{v_t} \frac{f_j(\mu_t) - f_0(\mu_t)}{(f_0(\mu_t) + 1)(f_j(\mu_t) + 1)} \\
 &= \frac{1-b}{v_t} \frac{C_f \sum_{i=k}^{\infty} w_i \cdot B_{[a_i, b_i]}(\mu_t)}{(f_0(\mu_t) + 1)(f_j(\mu_t) + 1)} \\
 &\leq \frac{1-b}{b} \cdot \frac{C_f \cdot 2w_k}{1 \times 1} \\
 &\leq 1 \cdot 2C_f w_k \\
 &\leq \frac{w_k}{3}
 \end{aligned}$$

Here the third last inequality comes from  $v_t \geq b$ ,  $f_0(v) \geq 0$ ,  $f_j(v) \geq 0$ , and the fact that

$$\sum_{i=k}^{\infty} w_i B_{[a_i, b_i]}(\mu_t) \leq \sum_{i=k}^{\infty} 3^{-i!} \cdot 1 \leq 3^{-k!} \sum_{i=0}^{\infty} 3^{-i} \leq \frac{2}{3} \cdot 3^{-k!} < 2w_k.$$

The second last inequality comes from  $b = \frac{6C_f + 1}{2} \geq \frac{1}{2}$ . The lastest inequality comes from the fact that  $6C_f < 1$ .

Now we propose a lemma:

**Lemma 16.** *For Bernoulli distributions  $Ber(p)$  and  $Ber(p + \epsilon)$  with  $\frac{1}{2} \leq p \leq p + \epsilon \leq \frac{1}{2} + C$ , we have*

$$KL(p||p + \epsilon) \leq \frac{1}{\ln 2} \frac{4}{1 - 4C^2} \epsilon^2.$$

*Proof.*

$$\begin{aligned}
 KL(p||p + \epsilon) &= p \log\left(\frac{p}{p + \epsilon}\right) + (1 - p) \log\left(\frac{1 - p}{1 - p - \epsilon}\right) \\
 &= \frac{1}{\ln 2} \cdot \left( p(-\ln(1 + \frac{\epsilon}{p})) + (1 - p) \ln(1 + \frac{\epsilon}{1 - p - \epsilon}) \right) \\
 &\leq \frac{1}{\ln 2} \cdot \left( p(-\frac{\epsilon}{p + \epsilon}) + (1 - p) \frac{\epsilon}{1 - p - \epsilon} \right) \\
 &= \frac{1}{\ln 2} \cdot \frac{\epsilon^2}{(p + \epsilon)(1 - p - \epsilon)} \\
 &\leq \frac{1}{\ln 2} \cdot \frac{\epsilon^2}{(\frac{1}{2} + C)(\frac{1}{2} - C)} \\
 &\leq \frac{1}{\ln 2} \cdot \frac{1}{\frac{1}{4} - C^2} \epsilon^2.
 \end{aligned}$$

Here the third line comes from the fact that  $\frac{v}{1+v} \leq \ln v \leq v$ . ■



Let us come back to the proof of the theorem. Since  $\frac{D_0(v_t)}{v_t} \geq b \geq \frac{1}{2}$ , and the fact that

$$\begin{aligned} \frac{D_j(v_t)}{v_t} &\leq \frac{D_j(b + \frac{1-b}{3})}{b + \frac{1-b}{3}} \\ &= 3 \cdot \frac{b + (1-b)(1 - G_j(\frac{1}{3}))}{1 + 2b} \\ &= 3 \frac{b + (1-b) \frac{f_j(\frac{1}{3})}{f_j(\frac{1}{3})+1}}{1 + 2b} \\ &= 3 \frac{b + (1-b) \frac{C_f}{C_f+1}}{1 + 2b}. \end{aligned}$$

The first inequality is because  $\frac{D(v)}{v}$  is non-increasing and the fact that  $v_t \geq b + \frac{1-b}{3}$  if  $v_t \in [a_k, b_k], k \geq 1$ . The last equality comes from the fact that  $f_j(\frac{1}{3}) = C_f$ . Now we specify the constants: let  $C_f = \frac{1}{60}, b = \frac{6C_f+1}{2} = \frac{11}{20}$ . Plug in these constant values and we get:

$$\frac{D_j(v_t)}{v_t} \leq \frac{340}{427} < \frac{5}{6}.$$

According to Lemma 16, we have:

$$KL\left(\frac{D_0(v_t)}{v_t} \parallel \frac{D_j(v_t)}{v_t}\right) \leq \frac{1}{\ln 2} \cdot \frac{4}{1 - 4 \cdot (\frac{5}{6} - \frac{1}{2})^2} \cdot \left(\frac{w_k}{3}\right)^2 = \frac{1}{\ln 2} \cdot \frac{36}{5} \cdot \frac{w_k^2}{9} = \frac{1}{\ln 2} \cdot \frac{4w_k^2}{5}.$$

. Recall that

$$\begin{aligned} &KL(\mathbb{P}_0(\mathbf{r}_n) \parallel \mathbb{P}_j(\mathbf{r}_n)) \\ &= \sum_{t=1}^t \mathbb{P}_0\left(\frac{v_t - b}{1 - b} \in I_j\right) \cdot KL\left(\frac{D_0(v_t)}{v_t} \parallel \frac{D_j(v_t)}{v_t}\right) \\ &\leq \frac{1}{\ln 2} \cdot \frac{4}{5} w_k^2 \cdot \sum_{t=1}^n \mathbb{P}_0[\mu_t \in I_j]. \end{aligned}$$

Therefore, we have:

$$\begin{aligned} &\mathbb{E}_{\mathbb{P}_j}[h(\mathbf{r}_n)] - \mathbb{E}_{\mathbb{P}_0}[h(\mathbf{r}_n)] \\ &\leq \frac{M}{2} \sqrt{2 \ln 2 \cdot \frac{1}{\ln 2} \cdot \frac{4}{5} w_k^2 \cdot \sum_{t=1}^n \mathbb{P}_0[\mu_t \in I_j]} \\ &\leq \frac{4M \cdot w_k}{5} \cdot \sqrt{\sum_{t=1}^n \mathbb{P}_0[\mu_t \in I_j]}. \end{aligned}$$

Now, let  $h(\mathbf{r}_n)$  be  $N_j = |\{t | \mu_t \in I_j, t = 1, 2, \dots, n_k\}|$ , and we know that  $M = n_k$ . Since  $n_k = \lceil \frac{1}{k} \frac{w_{k-1}}{w_k^3} \rceil$ , we conduct the pricing for  $n_k$  times and have:

$$\mathbb{E}_{\mathbb{P}_j}[N_j] - \mathbb{E}_{\mathbb{P}_0}[N_j] \leq \frac{4M \cdot w_k}{5} \cdot \sqrt{\sum_{t=1}^n \mathbb{P}_0[\mu_t \in I_j]} = \frac{4M \cdot w_k}{5} \cdot \sqrt{\mathbb{E}_{\mathbb{P}_0}[N_j]}.$$

Sum over  $j = 1, 2, \dots, Q_k$  of the inequality above, and we take an average to get:

$$\begin{aligned}
 \frac{1}{Q_k} \cdot \sum_{j=1}^{Q_k} \mathbb{E}_{\mathbb{P}_j}[N_j] &\leq \frac{1}{Q_k} \sum_{j=1}^{Q_k} \mathbb{E}_{\mathbb{P}_0}[N_j] + \frac{1}{Q_k} \frac{4}{5} n_k \cdot w_k \sum_{j=1}^{Q_k} \sqrt{\mathbb{E}_{\mathbb{P}_0}[N_j]} \\
 &= \frac{1}{Q_k} \cdot n_k + \frac{1}{Q_k} \frac{4}{5} n_k \cdot w_k \sum_{j=1}^{Q_k} \sqrt{\mathbb{E}_{\mathbb{P}_0}[N_j]} \\
 &\leq \frac{n_k}{Q_k} + \frac{4}{5} \frac{n_k}{Q_k} \cdot w_k \cdot \sqrt{Q_k \cdot \sum_{j=1}^{Q_k} \mathbb{E}_{\mathbb{P}_0}[N_j]} \\
 &= \frac{n_k}{Q_k} + \frac{4}{5} \frac{n_k}{Q_k} \cdot w_k \cdot \sqrt{Q_k n_k} \\
 &\leq \frac{3}{k} \cdot \frac{4}{5} \frac{3}{k} \frac{1}{w_k^2} \sqrt{\frac{3}{k^2} \cdot \frac{w_{k-1}^2}{w_k^4}} \\
 &= \frac{3}{k} \frac{1}{w_k^2} + \frac{4\sqrt{3}}{5\sqrt{k}} \frac{1}{k} \frac{w_{k-1}}{w_k^3} \\
 &\leq 0.9 \cdot n_k, \text{ for } k \geq 3.
 \end{aligned} \tag{11}$$

In Equation 11, the first line comes from the summation; the second (and the fourth) line is because  $\sum_{j=1}^{Q_k} \mathbb{E}_{\mathbb{P}_0}[N_j] = \mathbb{E}_{\mathbb{P}_0}[\sum_{j=1}^{Q_k} N_j] = n_k$ ; the third line is an application of Cauchy-Schwartz's Inequality; the fifth line is derived by plugging in  $Q_k = \frac{w_{k-1}}{3w_k}, n_k \leq \frac{1}{k} \cdot \frac{w_{k-1}}{w_k^3}, w_k = 3^{-k!}$ ; the last line is just calculations. Therefore, under distribution  $d_j$ , the policy  $\pi$  is expected to choose an  $v_t \notin I_j$  for at least  $0.1n_k$  times, which will bring a regret  $0.1n_k \cdot C_j \cdot w_k = \frac{1}{600} n_k \cdot w_k = \frac{1}{k} \cdot w_k^{\frac{1}{k}-2}$ . Since  $n_k = \frac{1}{k} \cdot w_k^{\frac{1}{k}-3}$ , we know that  $Regret = \Omega((n_k)^{\frac{2}{3}-\frac{1}{3k}})$  up to logarithmic factors. Therefore, we claim that for any  $\delta > 0$ , no policy can achieve  $o(n_k^{\frac{2}{3}-\delta})$  for sufficiently large  $k$ .

This ends the proof of Theorem 7.

## C More Discussions

In this work, we have developed two ‘‘linear’’ approaches toward the agnostic dynamic pricing problem. There are, however, still some issues that we have a handful of insights instead of rigorous proof or empirical evidence. Here we would like to present these ideas that might serve as heuristics for further research.

### C.1 Differences between LP and LV

As we stated in Section 1, LP models our strategy while LV modes the nature. Also, a good (no-regret) LP algorithm approaches the best linear policy in total while a good LV algorithm approaches the global optimal price at each round. When we adopt a LP problem model, we indeed have very little information about the market valuation other than obvious features of the product to sell. In this situation, a linear pricing policy is tractable and transparent to the customers, but it is not guaranteed to present or approach the best price. When we adopt a LV problem model, it is assumed that we have already known all features of the selling session (not limited to the product itself), and the fluctuation caused by the market is independent to the product. Therefore, we may learn from the feature-pricing-feedback data over the time and estimate the noise distribution, which would help approaching the best price combining with a greedy policy. Here is a concrete example regarding vehicle owners, dealers and buyers that illustrates the difference between LP and LV:

In Session 1, suppose we are the owner and would like to sell our used car to a buyer/dealer. A 3rd-party evaluator will evaluate your car based on a few (but not all) factors, e.g., mileage, duration, condition and accident records, and then subtract a certain amount from the selling price of an identical new car. This amount is usually linearly or near-linearly dependent on these factors listed above. Remember that this selling price is proposed by we owners. In other words, we are the seller in this session, and the buyer/dealer would respond by accepting or declining the price we propose. Here we adopt a linear pricing policy because we do not have full information of the selling session, and therefore customers' valuation model is indeed unclear to us.

In Session 2, suppose we are the dealer and would like to sell a used car to a buyer. Car dealers usually have sufficient information on the vehicle and the market supply-demand relationship. At least, we know clearly about which features are related to customers’ valuations. Therefore, it is reasonable for us to assume a parametric noisy valuation model (possibly a LV model) on their customers, and we would optimize these parameters based on historical selling records. With the model being well-learned, we may approach the global optimal price every time. That we directly make assumptions on customers’ valuation model is reasonable since we dealers have sufficient information, but this could still be risky if the features we can observe are limited.

### C.2 Applying LP Algorithm to LV Problem Model

LP and LV’s are two distinctly different problems: the optimal prices in an LV problem is not necessarily linear w.r.t.  $x_t$ : when  $x_t = 0$  with a zero-mean noise, the expected reward of a price that is slightly larger than 0 would be positive while the expected reward of 0 price is exactly 0. Therefore, we believe that an optimal linear policy would suffer an  $\Omega(T)$  regret in some LV settings even with known noise distributions. However, if the noise distribution is parametric by some parameter  $\eta \in \mathbb{R}^k$ , then we might have a “pseudo-linear” policy  $\tilde{\eta} : [\theta_1, \theta_2, \dots, \theta_d, \eta_1, \eta_2, \dots, \eta_k]^\top$  that takes  $\tilde{x} := [x_1, x_2, \dots, x_d, \phi_1(u), \phi_2(u), \dots, \phi_k(u)]^\top$  as input, and outputs  $v = \tilde{\eta}^\top \tilde{x}$ . A similar linearization idea in non-feature pricing has been adopted in Wang et al. (2021) and achieves optimal regrets. However, it is still unknown whether their methods can be applied to this feature-based LV problem. The key to this approach is to figure out a (nearly-)linear action-to-reward mapping, but this seems really hard in this setting. Again, in an LV problem it is the valuation instead of the optimal price that is linear.

### C.3 The Hardness of Pricing versus Bandits

The generic feature-based dynamic pricing problem can be reduced to a contextual bandit problem with continuum action and infinite policy spaces, despite some literature that assumes a different acceptance/declination reward scope (see Bartók et al. (2014)). Therefore, the gap between a dynamic pricing problem and an ordinary (discrete-action and finite-policy) contextual bandit problem can be observed from three perspectives. Firstly, the pricing feedback contains more information than a bandit feedback: if  $v_t$  is accepted, then any  $v \leq v_t$  would have been accepted if it were proposed. We call this a “half-space information”. Secondly, a discrete action space might not contain the optimal or any near-optimal price that matches the minimax regret: the revenue curve can vary drastically with respect to the price (e.g., consider a noise whose pdf is a rescaled Weierstrass function). Thirdly, a finite policy space might not contain the global optimal or any near-optimal policy, either. This is possible even for a parametric policy space where the parameter space is infinite. Therefore, we cannot directly adopt the regret bounds of contextual bandits onto feature-based dynamic pricing problems unless there exists a rigorous reduction.

However, we notice that the three perspectives above are pointing at different directions: the “half-space information” makes pricing easier than bandits, while the other two discretization issues makes it harder. In fact, we might partially offset the “continuum action” issue with the “half-space information” just like what we did in this paper: the revenue curve is actually “half Lipschitz” that  $g(v_1, u, F) - g(v_2, u, F) \leq v_1 - v_2$  if  $v_1 \geq v_2$ . This helps our algorithms get rid of the Lipschitz assumption. However, this is not rich enough to substantially reduce the regret as we still use bandit algorithms to achieve a minimax rate in an LP problem, where the lower bound holds even for Lipschitz revenue curve. Therefore, a very important question occurs to us: what else could a pricing feedback provide other than the “half Lipschitz”? Technically speaking, does a pricing feedback contain high-order information of the revenue curve? Besides, remember that we still do not have a unified approach toward a finite near-optimal policy set. In this work, we discretize the noise distribution by  $\frac{1}{\gamma}$  grids, which indeed increases the regret bound. For more sophisticated feature-valuation mapping (e.g., a non-linear valuation model) that is hard to parameterize, maybe it is not suitable to just apply naive discretization methods.

As a result, pricing problem seems at least as hard as bandits, and it is still unclear whether or not we could completely solve the feature-based dynamic pricing via contextual bandit methods (even though the major contributions toward single-product dynamic pricing are from multi-armed-bandits-related approaches).

### C.4 Social Impacts

In this work, we mainly focus on an online-fashion pricing problem where only one product is sold to one customer at each round (time spot). Therefore, it is not likely to commit a pricing discrimination according to its rigorous definition (since the price fluctuation over time should not be treated as discrimination). However, there exist chances that our algorithm could be misused. Notice that each item is characterized by a feature vector  $x_t$ , which

might be used to capture more information, e.g., customers' behaviors. On the one hand, it is indeed a price discrimination if we propose differently-generated prices to customers with different personal features even at different time point as long as the market has not changed substantially. On the other hand, this would lead to a potential leakage of personal privacy. It is usually forbidden to collect and use personal information for commercial use, but the sellers would at least know what the customers have bought and how much they have paid. Even though the feature  $x_t$  can be encoded with cryptographic techniques such that it is still suitable for learning (e.g., a "fully-homomorphic encryption", or FHE), at least the proposed prices are informative and might reveal the customer's behaviors. Indeed, auctions are a method to avoid any pricing discrimination, but it is not practical in most of the situations happening in our daily life.