
Improved Regret and Contextual Linear Extension for Pandora’s Box and Prophet Inequality

Junyan Liu*

University of Washington
junyanl1@cs.washington.edu

Ziyun Chen*

University of Washington
ziyuncc@cs.washington.edu

Kun Wang*

Purdue University
wang5675@purdue.edu

Haipeng Luo

University of Southern California
haipengl@usc.edu

Lillian J. Ratliff

University of Washington
ratliff1@uw.edu

Abstract

We study the Pandora’s Box problem in an online learning setting with semi-bandit feedback. In each round, the learner sequentially pays to open up to n boxes with unknown reward distributions, observes rewards upon opening, and decides when to stop. The utility of the learner is the maximum observed reward minus the cumulative cost of opened boxes, and the goal is to minimize regret defined as the gap between the cumulative expected utility and that of the optimal policy. We propose a new algorithm that achieves $\tilde{O}(\sqrt{nT})$ regret after T rounds, which improves the $\tilde{O}(n\sqrt{T})$ bound of Agarwal et al. [2024] and matches the known lower bound up to logarithmic factors. To better capture real-life applications, we then extend our results to a natural but challenging contextual linear setting, where each box’s expected reward is linear in some known but time-varying d -dimensional context and the noise distribution is fixed over time. We design an algorithm that learns both the linear function and the noise distributions, achieving $\tilde{O}(nd\sqrt{T})$ regret. Finally, we show that our techniques also apply to the online Prophet Inequality problem, where the learner must decide immediately whether or not to accept a revealed reward. In both non-contextual and contextual settings, our approach achieves similar improvements and regret bounds.

1 Introduction

Pandora’s Box is a fundamental problem in stochastic optimization, initiated by Weitzman [1978]. In this problem, the learner is given n boxes, each associated with a known cost and a known reward distribution. The learner may open a box at a time, paying the corresponding cost to observe a realized reward. Based on the known distributions, the learner designs a policy that specifies the order of inspecting boxes and a stopping rule for when to halt the process and select a reward. Weitzman [1978] shows that the optimal policy computes a threshold (i.e., reservation value) for each box based on its cost and distribution, and opens boxes in descending order of these thresholds until a stopping condition is met. This model abstracts a wide range of sequential decision-making scenarios where

*Equal contribution

information is costly to acquire and decisions must be made under uncertainty. For example, in the hiring process, an employer interviews candidates one by one, where each interview incurs a cost (e.g., time), and each candidate’s performance is drawn from a known distribution (e.g., candidate’s profile). The employer must decide when to stop interviewing and make an offer, balancing the expected benefit of finding a better candidate in the future against the accumulated cost of continued search. Other real-world applications include online shopping and path planning [Atsidakou et al., 2024].

While the classic Pandora’s Box problem and its variants have been extensively studied (see survey [Beyhaghi and Cai, 2024]), recent work [Fu and Lin, 2020, Guo et al., 2021, Garmiry et al., 2024, Agarwal et al., 2024, Atsidakou et al., 2024] has started to explore settings where the reward distributions are *unknown*. This line of work is motivated by practical scenarios where distributions are unavailable upfront and must be learned through interaction. These formulations move beyond the classical setting by requiring the learner to adapt based on observed outcomes. Prior studies have considered probably approximately correct (PAC) guarantees, establishing the near-optimal sample complexity [Fu and Lin [2020], Guo et al. [2021]], and regret minimization over a T -round repeated game [Gergatsouli and Tzamos [2022], Garmiry et al. [2024], Agarwal et al. [2024], Atsidakou et al. [2024]]. In the basic regret minimization setting, the learner repeatedly solves a Pandora’s Box problem at each round and observes the rewards of the opened boxes (a.k.a. semi-bandit feedback), drawn from *fixed but unknown* distributions. However, even in this setting, a gap remains between the $\Omega(\sqrt{nT})$ lower bound of Garmiry et al. [2024] and the best known upper bound of $\tilde{O}(Un\sqrt{T})$ by Agarwal et al. [2024], where U is the magnitude of the utility function that can be as large as n .

In addition to the fixed-distribution setting, prior work studied the online Pandora’s Box problems with *time-varying* reward distributions, which better capture many real-world scenarios [Gergatsouli and Tzamos [2022], Atsidakou et al. [2024]]. For example, in the hiring process, the distribution of candidate quality may shift over time due to market trends or seasonal patterns. To model such dynamics, [Gergatsouli and Tzamos [2022]] consider a setting where the rewards are chosen by an oblivious adversary and show that sublinear regret is achievable when the algorithm competes against a benchmark weaker than the optimal policy, obtained by imposing constraints on the set of boxes it can open. Later, Garmiry et al. [2024] showed that in the same setting, when the algorithm competes against the optimal policy, no algorithm can achieve sublinear regret, even with full information feedback (i.e., the learner observes the rewards of all boxes regardless of which policy is played). To make the problem tractable while still allowing distributions to change over time, Atsidakou et al. [2024] propose a contextual model where, in each round, the learner observes a context vector for each box, whose optimal threshold can be approximated by a function of the context and an unknown vector. Their approach reduces the problem to linear-quadratic online regression, which provides generality but leads to a regret bound with poor dependence on the horizon T . In particular, if the function is linear, they achieve $\tilde{O}(nT^{5/6})$ regret bound with semi-bandit feedback (referred to as bandit feedback in their work).

Motivated by these limitations, we address two main challenges in this paper. First, we close the gap in the regret dependence on the number of boxes n . Second, we propose a new contextual model that allows the reward distributions to vary over time, and under this setting, we achieve improved dependence on the time horizon T . Our contributions are summarized as follows, and further related work is discussed in Appendix B.

- For the online Pandora’s Box problem, we propose a new algorithm that achieves a regret bound of $\tilde{O}(\sqrt{nT})$, improving upon the $\tilde{O}(Un\sqrt{T})$ bound by Agarwal et al. [2024] and matching the known $\Omega(\sqrt{nT})$ lower bound (up to logarithmic factors) [Garmiry et al., 2024]. Our algorithm builds on the optimism-based framework of Agarwal et al. [2024], but introduces a simple yet crucial modification. Specifically, we first construct empirical distributions for each box using the observed reward samples, and then shift probability mass to form optimistic distributions that encourage exploration. In contrast to the fixed-mass shifting scheme used in Agarwal et al. [2024], we adaptively reallocate mass at each point based on problem-dependent factors. As we demonstrate below, this adaptive construction is central to enabling a more refined regret analysis.
- We extend these results to the contextual linear setting, where each box’s expected reward is linear in some known but time-varying d -dimensional context and the noise distribution is fixed over time. Our algorithm preserves the overall structure of the non-contextual case, but modifies the construction of optimistic distributions to account for context-dependent variation. Since the reward

samples are no longer identically distributed, we adjust not only the probability mass but also the observed rewards themselves in a *value-optimistic* manner to maintain optimism. While the algorithm of [Atsidakou et al. \[2024\]](#) applies to our setting with $\tilde{O}(nT^{5/6})$ regret (see [Appendix B](#)), we establish $\tilde{O}(nd\sqrt{T})$ regret. As long as $d = o(T^{1/3})$, our algorithm yields a better regret guarantee.

- Finally, we extend our algorithms and analytical techniques to the Prophet Inequality problem in both non-contextual and contextual settings. In the non-contextual case, our algorithm achieves $\tilde{O}(\sqrt{nT})$ regret bound, which again improves upon the $\tilde{O}(n\sqrt{T})$ regret bound of [Agarwal et al. \[2024\]](#). For the contextual linear setup, we also establish a $\tilde{O}(nd\sqrt{T})$ regret bound. Due to space constraints, all results and proofs for the Prophet Inequality problem are deferred to the appendix.

Technical Overview for $\tilde{O}(\sqrt{nT})$ Improvement. Inspired by [Agarwal et al. \[2024\]](#), we decompose the regret at any round t by $\sum_i \text{Term}_{t,i}$, where $\text{Term}_{t,i}$ captures the difference in expected reward when replacing the optimistic distribution of the i -th box, denoted by $\hat{\mathcal{E}}_{t,i}$, with its true distribution D_i , while keeping all other boxes fixed. [Agarwal et al. \[2024\]](#) bound $\text{Term}_{t,i}$ by $\mathcal{O}(\text{TV}(\hat{\mathcal{E}}_{t,i}, D_i))$, where TV denotes the total variation distance. However, simply bounding $\text{Term}_{t,i}$ by the TV distance results in $\tilde{O}(n\sqrt{T})$ regret bound. To obtain a tighter bound, we take a different approach to handle $\text{Term}_{t,i}$. Let $\tilde{R}_i(\sigma_t, z)$ be the expected utility when using threshold vector σ_t , and the i -th box has a fixed value z . We then write $\text{Term}_{t,i} = \int_0^1 (F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz$ where F_D is the cumulative density function of distribution D . Our main techniques to bound this term are as follows.

Technique 1: Bernstein-type Bound. Both our analysis and algorithm design are based on a Bernstein-type concentration bound. In particular, our algorithm adaptively adjusts the probability mass when constructing empirical distributions, guided by this bound. This introduces an extra factor $\sqrt{F_{D_i}(\cdot)(1 - F_{D_i}(\cdot))}$ in the upper bound on $|F_{D_i}(\cdot) - F_{\mathcal{E}_{t,i}}(\cdot)|$, where $\mathcal{E}_{t,i}$ is the *empirical distribution* constructed at round t by assigning equal probability mass to iid samples from D_i .

Technique 2: Sharp Bound of $\frac{\partial}{\partial z} \tilde{R}_i(\sigma_t, z)$. We show 1-Lipschitz and monotonically-increasing properties for $\tilde{R}_i(\sigma_t, z)$, which imply that $0 \leq \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t, z) \leq 1$. Moreover, when z is small, we prove a sharper bound where $\frac{\partial}{\partial z} \tilde{R}_i(\sigma_t, z)$ is bounded by the probability that the first $i - 1$ boxes have values no larger than z .

Next, we apply the techniques above to obtain a sharper bound for $\text{Term}_{t,i} = \int_0^1 (F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz$. The high-level idea is as follows: when z is large, the Bernstein-type bound for $(F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z))$ contributes to a factor $\sqrt{F_{D_i}(z)(1 - F_{D_i}(z))}$, which is small since a large z implies a small $1 - F_{D_i}(z)$; on the other hand, when z is small, we use the sharper bound on $\frac{\partial}{\partial z} \tilde{R}_i(\sigma_t, z)$ to obtain a tighter bound.

Techniques for Contextual Linear Case. We first build a *virtual empirical distribution* $\hat{D}_{t,i}$ for each box i at round t , by using all available noise samples from this box plus the mean of the current round. This is virtual since we only observe non-i.i.d. reward samples, not the raw noises. To overcome this issue, we introduce a value-optimistic empirical distribution $\hat{\mathcal{E}}_{t,i}$ by using the linear mean model to shift the observed samples to overestimate $\hat{D}_{t,i}$. We then decompose $\text{Term}_{t,i}$ into three terms, two of which can be handled in a similar way as in the non-contextual case. The remaining piece is $\mathbb{E}_{z \sim \mathcal{E}_{t,i}}[\tilde{R}_i(\sigma_{\hat{\mathcal{E}}_t}; z)] - \mathbb{E}_{z \sim \hat{D}_{t,i}}[\tilde{R}_i(\sigma_{\hat{\mathcal{E}}_t}; z)]$ which measures the mean estimation error. This term is bounded by applying the 1-Lipschitz property of $\tilde{R}_i(\sigma_{\hat{\mathcal{E}}_t}; z)$ and contextual linear bandit techniques.

2 Preliminaries

In this section, we introduce online Pandora’s Box problem in both non-contextual and contextual settings. Due to the limited space, the Prophet Inequality problem is deferred to [Appendix A](#).

Online Pandora’s Box. In this problem, the learner is given n boxes, each with a known cost $c_i \in [0, 1]$ and an unknown fixed reward distribution D_i supported on $[0, 1]$. The learner plays a T -round repeated game where at each round $t \in [T]$, the learner decides an order to inspect boxes

Algorithm 1 Generic descending threshold algorithm for Pandora’s Box

Input: threshold $\sigma = (\sigma_1, \dots, \sigma_n)$.

Initialize: $V_{\max} = -\infty$. Let $\pi_i \in [n]$ be the box with the i -th largest threshold.

The environment generates an unknown realization $v = (v_1, \dots, v_n)$.

for $i = 1, \dots, n$ **do**

 Pay c_{π_i} to open box π_i , observe v_{π_i} , and update $V_{\max} = \max\{V_{\max}, v_{\pi_i}\}$.

 If $i = n$ or $V_{\max} \geq \sigma_{\pi_{i+1}}$, then stop and take reward $V_{\max}$.

and a stopping rule. Upon opening a box i , the learner pays cost c_i and receives a reward $v_{t,i} \in [0, 1]$ independently sampled from D_i . Once the stopping rule is met, she selects the highest observed reward so far. The utility for the round is the chosen reward minus the total cost of opened boxes.

Contextual Linear Pandora’s Box. This setting extends the above by allowing each box’s distribution to vary across rounds. Specifically, each box $i \in [n]$ is associated with a known cost $c_i \in [0, 1]$, an unknown parameter $\theta_i \in \mathbb{R}^d$, and a fixed unknown noise distribution D_i with zero-mean, supported on $[-\frac{1}{4}, \frac{1}{4}]$. At the beginning of each round $t \in [T]$, the learner observes a context $x_{t,i} \in \mathbb{R}^d$ for each box $i \in [n]$. These contexts $\{x_{t,i}\}_{i \in [n]}$ can be chosen arbitrarily by an adaptive adversary. Based on these contexts and past observations, the learner then decides an inspection order and a stopping rule to open boxes. Different from the non-contextual setup, upon opening box i , the learner observes the reward $v_{t,i} = \eta_{t,i} + \mu_{t,i}$ where $\mu_{t,i} = \theta_i^\top x_{t,i} \in [\frac{1}{4}, \frac{3}{4}]$ is the mean and $\eta_{t,i}$ is a noise independently drawn from D_i .² That is, the reward is drawn from the $[0, 1]$ -bounded distribution $D_{t,i} = \mu_{t,i} + D_i$ which shifts D_i by the context-dependent mean $\mu_{t,i}$. Following standard assumptions in contextual linear bandits, we assume that $\|\theta_i\|_2 \leq 1$ and $\|x_{t,i}\|_2 \leq 1$ for all t, i .

Optimal Policy and Regret. To measure the performance, we compete our algorithm against the per-round optimal policy. For both settings, the optimal policy at round t is the Weitzman’s algorithm [Weitzman, 1978] which operates with the distribution $D_t = (D_{t,1}, \dots, D_{t,n})$. In the non-contextual setting, the distribution is fixed across rounds i.e., $D_t = D = (D_1, \dots, D_n)$, and thus, we will retain the notation D_t in the following. More specifically, Weitzman’s algorithm proceeds by computing a threshold $\sigma_{t,i}^*$ for each box i , which is the solution of $\mathbb{E}_{x \sim D_{t,i}}[(x - \sigma_{t,i}^*)_+] = c_i$ where $(x)_+ = \max\{0, x\}$. Then, the algorithm inspects the boxes in descending order of $\sigma_{t,i}^*$ and stops as soon as the maximum observed reward so far exceeds the threshold of the next unopened box.

To formally define regret, we introduce the following notations. For any product distribution $D = (D_1, \dots, D_n)$, let $\sigma_D = (\sigma_{D,1}, \dots, \sigma_{D,n})$ denote the optimal threshold vector, where each $\sigma_{D,i}$ is the solution of $\mathbb{E}_{x \sim D_i}[(x - \sigma_{D,i})_+] = c_i$. For any threshold vector $\sigma = (\sigma_1, \dots, \sigma_n) \in [0, 1]^n$ and reward realization $v = (v_1, \dots, v_n) \in [0, 1]^n$, we use $W(\sigma; v) \subseteq [n]$ to denote boxes opened by Algorithm 1. Further, we define the utility function and expected utility function respectively as:

$$R(\sigma; v) = \max_{i \in W(\sigma; v)} v_i - \sum_{i \in W(\sigma; v)} c_i, \quad \text{and} \quad R(\sigma; D) = \mathbb{E}_{v \sim D}[R(\sigma; v)].$$

With these definitions, the optimal expected utility is $R(\sigma_D; D_t)$ since when $v \sim D_t$, the set $W(\sigma_D; v)$ corresponds to the boxes opened by the Weitzman’s algorithm. Let $\sigma_t \in \mathbb{R}^n$ denote the threshold vector selected by the algorithm at round t . Throughout the paper, all our algorithms follow Algorithm 1 using σ_t as input at round t , and thus the expected utility our algorithm at round t is $R(\sigma_t; D_t)$. The cumulative regret in T rounds Reg_T is then defined as follows: $\text{Reg}_T = \mathbb{E} \left[\sum_{t=1}^T R(\sigma_D; D_t) - R(\sigma_t; D_t) \right]$, where $D_t = D$ for the non-contextual case.

3 $\tilde{\mathcal{O}}(\sqrt{nT})$ Regret Bound for Non-Contextual Pandora’s Box

In this section, we first present Algorithm 2 for the non-contextual Pandora’s Box problem and the main result, and then discuss the analysis.

²The boundedness assumptions on the noise and mean ensure that rewards lie in $[0, 1]$, which avoids re-deriving standard results under a different scaling. Our results readily extends to sub-Gaussian rewards. We refer readers to Appendix C.7 for details.

3.1 Algorithm and Main Result

While [Algorithm 2](#) is inspired by the principle of *optimism in the face of uncertainty*, its application differs from that in the classic multi-armed bandit setting, such as in the Upper Confidence Bound (UCB) algorithm [[Auer et al., 2002](#)]. In bandit problems, optimism is typically applied to the mean reward of each arm. In contrast, the Pandora’s Box problem requires learning the entire distribution of each box, particularly its cumulative distribution function (CDF). We follow a similar idea to [Agarwal et al. \[2024\]](#), but employ a different approach in the construction of an optimistic estimate of the distribution. The threshold σ_t is computed from this optimistic distribution, which implicitly ensures optimism in the algorithm’s decisions. More formally, we begin by introducing the notion of stochastic dominance.

Definition 3.1.1 (Stochastic dominance). Let D and E be two probability distributions with CDFs F_D and F_E , respectively. If $\mathbb{P}_{X \sim E}(X \geq a) \geq \mathbb{P}_{Y \sim D}(Y \geq a)$ for all $a \in \mathbb{R}$, we say that E stochastically dominates D , and we denote this by $E \succeq_{SD} D$.

For any two product distributions $D = (D_1, \dots, D_n)$ and $\mathcal{E} = (\mathcal{E}_1, \dots, \mathcal{E}_n)$, we use $\mathcal{E} \succeq_{SD} D$ to indicate that $\mathcal{E}_i \succeq_{SD} D_i$ for all $i \in [n]$. Let $m_{t,i}$ be the number of samples for each box i at round t and let $t_i(j)$ be the round that the j -th reward sample of box i is observed by the learner. We define empirical distribution and optimistic distribution as follows.

- **Empirical distribution $\mathcal{E}_{t,i}$.** For each box $i \in [n]$, we use $m_{t,i}$ i.i.d. samples $\{v_{t_i(j),i}\}_{j=1}^{m_{t,i}}$ to construct $\mathcal{E}_{t,i}$ with respect to the underlying distribution D_i by assigning each sample with probability mass $\frac{1}{m_{t,i}}$. Specifically, $F_{\mathcal{E}_{t,i}}(x) = \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \mathbb{I}\{v_{t_i(j),i} \leq x\}$ for all $x \in [0, 1]$.
- **Optimistic distribution $\hat{\mathcal{E}}_{t,i}$.** Let $L = 4 \log(2nT^2/\delta)$. The CDF of $\hat{\mathcal{E}}_{t,i}$ is constructed as follows:

$$F_{\hat{\mathcal{E}}_{t,i}}(x) = \begin{cases} 1, & x = 1; \\ \max \left\{ 0, F_{\mathcal{E}_{t,i}}(x) - \sqrt{\frac{2F_{\mathcal{E}_{t,i}}(x)(1-F_{\mathcal{E}_{t,i}}(x))L}{m_{t,i}}} - \frac{L}{m_{t,i}} \right\}, & 0 \leq x < 1. \end{cases} \quad (1)$$

As verified in [Lemma C.1.4](#), the construction in [Eq. \(1\)](#) ensures that $F_{\hat{\mathcal{E}}_{t,i}}(\cdot)$ is a valid CDF. We denote $\mathcal{E}_t = (\mathcal{E}_{t,1}, \dots, \mathcal{E}_{t,n})$ and $\hat{\mathcal{E}}_t = (\hat{\mathcal{E}}_{t,1}, \dots, \hat{\mathcal{E}}_{t,n})$. This type of construction has been previously used in [[Guo et al., 2019, 2021](#)], but our use differs substantially in the analysis. The following lemma shows that the optimistic product distribution stochastically dominates the underlying product distribution.

Lemma 3.1.2. With probability at least $1 - \delta$, for all $t \in [T]$, we have $\hat{\mathcal{E}}_t \succeq_{SD} D$.

It is noteworthy that our approach to constructing the optimistic distribution is more adaptive than that of [[Agarwal et al., 2024](#)]. In particular, [Agarwal et al. \[2024\]](#) move a fixed amount of probability mass, approximately $1/\sqrt{m_{t,i}}$, to the maximal value that D_i can take. In contrast, we adjust the CDF based on a data-dependent confidence interval, resulting in a more fine-grained and adaptive optimistic distribution. As stated in the following theorem, this construction results in a tighter regret bound.

Theorem 3.1.3. For Pandora’s Box problem, with $\delta = T^{-1}$ [Algorithm 2](#) ensures $\text{Reg}_T = \tilde{O}(\sqrt{nT})$.

Again, our bound matches the $\Omega(\sqrt{nT})$ lower bound of [Gatmiry et al. \[2024\]](#) up to logarithmic factors. It is worth noting that their lower bound is proven under the easier full-information setting and thus applicable to our setting.

Remark 3.1.4 (High-probability regret bound). Under regret metric $\sum_{t=1}^T (R(\sigma_D; D) - R(\sigma_t; D))$, our proposed algorithm also achieves a high-probability bound of $\tilde{O}(\sqrt{nT})$ via two modifications only in our analysis. We refer readers to [Appendix C.6](#) for details.

3.2 Analysis

In this subsection, we sketch the main proof ideas for [Theorem 3.1.3](#) and summarize the new analytical techniques. We start by leveraging a monotonicity property for Pandora’s Box problem, established by [Guo et al. \[2021\]](#): for any distributions D, E , if $E \succeq_{SD} D$, then $R(\sigma_E; E) \geq R(\sigma_D; D)$. Since

Algorithm 2 Near-optimal algorithm for Pandora's Box

Input: confidence $\delta \in (0, 1)$, horizon T .

Initialize: open all boxes once and observe rewards. Set $m_{1,i} = 1$ for all $i \in [n]$.

for $t = 1, 2, \dots, T$ **do**

 For each $i \in [n]$, construct an optimistic distribution $\hat{\mathcal{E}}_{t,i}$ according to [Eq. \(1\)](#).

 Compute $\sigma_t = (\sigma_{t,1}, \dots, \sigma_{t,n})$ where $\sigma_{t,i}$ is the solution of $\mathbb{E}_{x \sim \hat{\mathcal{E}}_{t,i}}[(x - \sigma_{t,i})_+] = c_i$.

 Run [Algorithm 1](#) with σ_t to open a set of boxes, denoted by $\mathcal{B}_t$, and observe rewards $\{v_{t,i}\}_{i \in \mathcal{B}_t}$.

 Update counters $m_{t+1,i} = m_{t,i} + 1, \forall i \in \mathcal{B}_t$ and $m_{t+1,i} = m_{t,i}, \forall i \notin \mathcal{B}_t$.

the algorithm uses threshold $\sigma_t = \sigma_{\hat{\mathcal{E}}_t}$ at each round t and [Lemma 3.1.2](#) shows $\hat{\mathcal{E}}_t \succeq_{\text{SD}} D$, the regret at round t , denoted by $\text{Reg}(t)$, is bounded as

$$\text{Reg}(t) := \mathbb{E} \left[R(\sigma_D; D) - R(\sigma_t; \hat{\mathcal{E}}_t) + R(\sigma_t; \hat{\mathcal{E}}_t) - R(\sigma_t; D) \right] \leq \mathbb{E} \left[R(\sigma_t; \hat{\mathcal{E}}_t) - R(\sigma_t; D) \right].$$

For simplicity of the exposition, without loss of generality, we assume that $\sigma_{t,1} \geq \sigma_{t,2} \geq \dots \geq \sigma_{t,n}$. Following [Agarwal et al. \[2024\]](#), we define

$$\mathcal{H}_{t,0} = \hat{\mathcal{E}}_t, \quad \text{and} \quad \mathcal{H}_{t,i} = (D_1, \dots, D_i, \hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n}) \quad \text{for all } t, i. \quad (2)$$

Therefore $\mathcal{H}_{t,n} = D$ so that

$$R(\sigma_t; \hat{\mathcal{E}}_t) - R(\sigma_t; D) = \sum_{i=1}^n R(\sigma_t; \mathcal{H}_{t,i-1}) - R(\sigma_t; \mathcal{H}_{t,i}). \quad (3)$$

The decomposition in [Eq. \(3\)](#) breaks the reward difference into a sequence of steps, where each step replaces one coordinate of the optimistic distribution with its true counterpart, thereby allowing for isolation of the effect of each individual coordinate change.

Let $E_{\sigma,i}$ be the event that box i is opened under the threshold σ , and let $E_{\sigma,i}^c$ be the corresponding complementary event. We define $Q_{t,i}$ as the probability that σ_t opens box i given the true product distribution D . Using the same argument as in [\[Agarwal et al., 2024\]](#), we have that

$$\begin{aligned} & \mathbb{E} [R(\sigma_t; \mathcal{H}_{t,i-1}) - R(\sigma_t; \mathcal{H}_{t,i})] \\ &= \mathbb{E} \left[Q_{t,i} \underbrace{\left(\mathbb{E}_{x \sim \mathcal{H}_{t,i-1}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{y \sim \mathcal{H}_{t,i}} [R(\sigma_t; y) \mid E_{\sigma_t,i}] \right)}_{\text{Term}_{t,i}} \right]. \end{aligned} \quad (4)$$

The key step is to carefully bound $\text{Term}_{t,i}$, and this is where our analysis *fundamentally diverges* from [Agarwal et al. \[2024\]](#). Before presenting our analysis, we explain below why the previous analysis fails to achieve the $\tilde{O}(\sqrt{nT})$ regret bound.

$\tilde{O}(Un\sqrt{T})$ bound of [Agarwal et al. \[2024\]](#). In the analysis of [Agarwal et al. \[2024\]](#), if the product distribution D is discrete, then they bound $\text{Term}_{t,i} \leq \mathcal{O}(U \cdot \text{TV}(D_i, \hat{\mathcal{E}}_{t,i}))$ (cf. [\[Agarwal et al., 2024, Lemma 1.3\]](#)) where U is the upper bound of the absolute value of function R and could be as large as n . Therefore,

$$\text{Reg}_T = \sum_{t=1}^T \text{Reg}(t) \leq \sum_{t=1}^T \mathbb{E} \left[R(\sigma_t; \hat{\mathcal{E}}_t) - R(\sigma_t; D) \right] \leq U \sum_{t=1}^T \sum_{i=1}^n \mathcal{O} \left(\mathbb{E} \left[Q_{t,i} \text{TV}(D_i, \hat{\mathcal{E}}_{t,i}) \right] \right).$$

Since their algorithm moves probability mass by a fixed amount, approximately $\tilde{O}(1/\sqrt{m_{t,i}})$ for each distribution i , a simple calculation gives $\text{TV}(D_i, \hat{\mathcal{E}}_{t,i}) \leq \tilde{O}(1/\sqrt{m_{t,i}})$, which in turn implies that

$$\text{Reg}_T \leq \tilde{O} \left(U \cdot \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^n \frac{Q_{t,i}}{\sqrt{m_{t,i}}} \right] \right) \leq \tilde{O} \left(U \cdot \sqrt{\mathbb{E} \left[n \sum_{t=1}^T \sum_{i=1}^n Q_{t,i} \right]} \right) \leq \tilde{O} (Un\sqrt{T}), \quad (5)$$

where the second inequality follows from the standard analysis of optimistic algorithms (e.g., [\[Slivkins, 2019\]](#)), and the last inequality uses the fact that $\sum_{i=1}^n Q_{t,i} \leq n$ for all t . It is noteworthy that

$\sum_{i=1}^n Q_{t,i} = 1$ appears in multi-armed bandit problems since only one arm is played in each round. However, in the problems we consider, there is a chance to open all boxes in one round.

Refined analysis for $\tilde{\mathcal{O}}(\sqrt{nT})$ bound. To get the $\tilde{\mathcal{O}}(\sqrt{nT})$ regret bound, it suffices to show $\text{Reg}(t) \leq \tilde{\mathcal{O}}(\sqrt{\sum_i Q_{t,i}/m_{t,i}})$ since

$$\text{Reg}_T \leq \tilde{\mathcal{O}} \left(\mathbb{E} \left[\sum_{t=1}^T \sqrt{\sum_{i \in [n]} \frac{Q_{t,i}}{m_{t,i}}} \right] \right) \leq \tilde{\mathcal{O}} \left(\mathbb{E} \left[\sqrt{T \sum_{t=1}^T \sum_{i \in [n]} \frac{Q_{t,i}}{m_{t,i}}} \right] \right) \leq \tilde{\mathcal{O}}(\sqrt{nT}),$$

where the second inequality follows from the Cauchy–Schwarz inequality, and the last inequality results from the bound $\sum_{t,i} Q_{t,i}/m_{t,i} \leq \tilde{\mathcal{O}}(n)$. For any threshold $\sigma \in [0, 1]^n$ and $z \in [0, 1]$, define

$$\tilde{R}_i(\sigma; z) := \mathbb{E}_{x \sim \mathcal{H}_{t,i}} [R(\sigma; (x_1, \dots, x_{i-1}, z, x_{i+1}, \dots, x_n)) \mid E_{\sigma,i}]. \quad (6)$$

Based on these definitions and the fact that $E_{\sigma_t,i}$ is independent of i -th box's reward, we write

$$\text{Term}_{t,i} = \mathbb{E}_{z \sim \hat{\mathcal{E}}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim D_i} [\tilde{R}_i(\sigma_t; z)].$$

The following lemma is a key enabler for our refined analysis of both problems.

Lemma 3.2.1 (1-Lipschitzness and monotonicity). *Consider the Pandora's box problem. For all $t \in [T]$, $i \in [n]$, the function $\tilde{R}_i(\sigma_t; z)$ is 1-Lipschitz and monotonically-increasing with respect to z .*

Since $\tilde{R}_i(\sigma_t; z)$ is 1-Lipschitz with respect to z , the map $\tilde{R}_i(\sigma_t; z)$ is absolutely continuous on $[0, 1]$, which implies that it is differentiable almost everywhere in $[0, 1]$. By the fundamental theorem of calculus, for any $x \in [0, 1]$, we have that $\tilde{R}_i(\sigma_t; x) = \tilde{R}_i(\sigma_t; 0) + \int_0^x \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz$. Then it is straightforward to deduce that

$$\begin{aligned} \text{Term}_{t,i} &= \int_0^1 (1 - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz - \int_0^1 (1 - F_{D_i}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz \\ &= \int_0^1 (F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz \\ &\leq \underbrace{\int_{\sigma_{t,i+1}}^1 |F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| dz}_{A_{t,i}} + \underbrace{\int_0^{\sigma_{t,i+1}} (F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz}_{B_{t,i}}, \end{aligned}$$

where the inequality holds since $\tilde{R}_i(\sigma_t; z)$ is 1-Lipschitz with respect to z and thus its derivative is bounded in $[-1, 1]$ almost everywhere. Here, we decompose $\text{Term}_{t,i}$ into two parts $A_{t,i}$ and $B_{t,i}$ to isolate contributions from different value regions, enabling a problem-dependent analysis that yields a sharper regret bound. An advantage of this analysis is to avoid using [Agarwal et al., 2024, Lemma 1.3], which incurs a linear dependence on U . Indeed, we have

$$\begin{aligned} A_{t,i} &\leq \int_{\sigma_{t,i+1}}^1 |F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)| dz + \int_{\sigma_{t,i+1}}^1 |F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| dz \\ &\leq \mathcal{O} \left(\frac{L}{m_{t,i}} + \int_{\sigma_{t,i+1}}^1 \sqrt{\frac{F_{D_i}(z)(1 - F_{D_i}(z))L}{m_{t,i}}} dz + \int_{\sigma_{t,i+1}}^1 \sqrt{\frac{F_{\mathcal{E}_{t,i}}(z)(1 - F_{\mathcal{E}_{t,i}}(z))L}{m_{t,i}}} dz \right) \\ &\leq \mathcal{O} \left(\frac{L}{m_{t,i}} + \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} + \sqrt{\frac{(1 - F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1}))L}{m_{t,i}}} \right) \\ &\leq \mathcal{O} \left(\sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right), \end{aligned}$$

where the second inequality bounds $|F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)|$ by a Bernstein-type Dvoretzky-Kiefer-Wolfowitz inequality (cf. Lemma C.1.2 in Appendix C.1) and bounds $|F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)|$ by the

construction [Eq. \(1\)](#); the third inequality bounds $1 - F_{D_i}(z) \leq 1 - F_{D_i}(\sigma_{t,i+1})$, and an analogous inequality holds for $1 - F_{\mathcal{E}_{t,i}}(z)$ since $z \geq \sigma_{t,i+1}$; the last inequality replaces $F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1})$ with $F_{D_i}(\sigma_{t,i+1})$ by paying a term of the same order (cf. [Eq. \(19\)](#)).

To handle the term $B_{t,i}$, instead of bounding the derivative by one as before, we introduce the following lemma which more tightly bounds the derivative by the probability that the first $i - 1$ boxes have values no larger than z , conditioning on opening box i .

Lemma 3.2.2. *Suppose that $\sigma_{t,1} \geq \sigma_{t,2} \geq \dots \geq \sigma_{t,n}$. Pandora's Box problem satisfies*

$$0 \leq \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) \leq \prod_{j < i} F_{D_j}(z) / Q_{t,i}, \quad \forall z \in [0, \sigma_{t,i+1}).$$

For shorthand, define $a_i^z = \prod_{j \leq i} F_{D_j}(z)$. By [Lemma 3.2.2](#) and some calculations (deferred to [Appendix C](#)), we have that

$$\begin{aligned} B_{t,i} &\leq \int_0^{\sigma_{t,i+1}} |F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)| \frac{\sqrt{a_{i-1}^z} \sqrt{a_{i-1}^z - 1}}{Q_{t,i}} dz \\ &\leq \int_0^{\sigma_{t,i+1}} \left(|F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)| + |F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| \right) \sqrt{\frac{a_{i-1}^z}{Q_{t,i}}} dz \\ &\leq \mathcal{O} \left(\frac{L}{m_{t,i}} + \frac{1}{\sqrt{Q_{t,i}}} \int_0^1 \sqrt{a_{i-1}^z (1 - F_{D_i}(z))} \sqrt{\frac{L}{m_{t,i}}} dz \right), \end{aligned}$$

where the second inequality follows from $a_{i-1}^z \leq a_{i-1}^{\sigma_{t,i+1}} \leq a_{i-1}^{\sigma_{t,i}} = Q_{t,i}$ for all $z < \sigma_{t,i+1}$, and the last inequality repeats the same argument used to bound $A_{t,i}$ to bound $|F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)|$ and $|F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)|$, respectively.

Plugging the above bounds of $A_{t,i}$ into [Eq. \(3\)](#) and [Eq. \(4\)](#), the regret contribution of $\{A_{t,i}\}_i$ at round t (omitting $L/m_{t,i}$ and some constant) is bounded by

$$\sum_{i \in [n]} Q_{t,i} \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} \leq \sqrt{\sum_{i \in [n]} Q_{t,i} (1 - F_{D_i}(\sigma_{t,i+1}))} \sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}} \leq \sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}},$$

where the first inequality follows from the Cauchy-Schwarz inequality, and the second inequality follows from $\sum_{i \in [n]} Q_{t,i} (1 - F_{D_i}(\sigma_{t,i+1})) \leq 1$. Indeed, since $Q_{t,i} = a_{i-1}^{\sigma_{t,i}}$ and $\sigma_{t,1} \geq \sigma_{t,2} \dots \geq \sigma_{t,n}$, we have that $Q_{t,i} \leq a_{i-1}^{\sigma_{t,j+1}}$, which in turn yields the telescoping sum $\sum_{i=1}^n Q_{t,i} (1 - F_{D_i}(\sigma_{t,i+1})) \leq \sum_{i=1}^n (a_{i-1}^{\sigma_{t,j+1}} - a_i^{\sigma_{t,j+1}}) \leq 1$.

Similarly, the regret contribution of $\{B_{t,i}\}_i$ at round t (omitting $L/m_{t,i}$ and some constant) is bounded by

$$\sum_{i \in [n]} \int_0^1 \sqrt{a_{i-1}^z - a_i^z} \sqrt{\frac{LQ_{t,i}}{m_{t,i}}} dz \leq \int_0^1 \sqrt{\sum_{i \in [n]} (a_{i-1}^z - a_i^z)} \sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}} dz \leq \sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}},$$

where the first inequality uses the Cauchy-Schwarz inequality and the second inequality uses the telescoping sum and $a_i^z \leq 1$ for all i, z . Hence, the cumulative regret is bounded by $\tilde{\mathcal{O}}(\sqrt{\sum_i Q_{t,i}/m_{t,i}})$. Combining all the results completes the proof of [Theorem 3.1.3](#).

4 $\tilde{\mathcal{O}}(nd\sqrt{T})$ Regret Bound for Contextual Linear Pandora's Box

In this section, we present [Algorithm 5](#) (in [Appendix D](#)) for the contextual linear Pandora's Box problem. While our algorithm inherits the main structure of [Algorithm 2](#), it introduces key modifications to ensure optimism, which is critical in the contextual setting. To ensure the optimism, we need to construct an optimistic product distribution which stochastically dominates the underlying product distribution D_t for each round t . In the non-contextual case, the optimistic distribution is constructed by using *iid samples* to first construct empirical distribution and then move the probability mass.

However, due to context-dependent shifts, this strategy no longer applies. To address this, we leverage the fact that the shift in each box's distribution is determined by a fixed but unknown vector θ_i .

A natural idea here is to construct empirical distribution by using samples $\{z_{t_i(j),i}^t\}_{j=1}^{m_{t,i}}$ where $z_{t_i(j),i}^t = \eta_{t_i(j),i} + \mu_{t,i}$, since they can be treated as $m_{t,i}$ iid samples drawn from $D_{t,i}$. However, the learner never directly observes the noise sample $\eta_{t_i(j),i}$. To overcome this, we maintain an estimation $\hat{\theta}_{t,i}$ at each round t by using ridge regression on past observations. Specifically, for each opened box $i \in \mathcal{B}_t$, we make the following updates: $\hat{\theta}_{t,i} = V_{t,i}^{-1} \sum_{s \leq t: i \in \mathcal{B}_s} x_{s,i} v_{s,i}$ where $V_{t,i} = I + \sum_{s \leq t: i \in \mathcal{B}_s} x_{s,i} x_{s,i}^\top$. Let $\mathcal{X}$ denote the context space. By the analysis of linear bandits (e.g., [Abbasi-Yadkori et al., 2011]), for each box $i \in [n]$, there exists a *known* function $\beta_i^t : \mathcal{X} \times (0, 1) \rightarrow \mathbb{R}_{>0}$ such that with probability at least $1 - \delta$, for all $x \in \mathcal{X}$ and all t ,

$$\left| x^\top (\theta_i - \hat{\theta}_{t-1,i}) \right| \leq \beta_i^t(x, \delta), \text{ where } \beta_i^t(x, \delta) = \mathcal{O} \left(\|x\|_{V_{t-1,i}^{-1}} \sqrt{\log(n\delta^{-1}) + d \log(1 + T/d)} \right).$$

With a confidence radius of θ , we construct a *value-optimistic* empirical distribution $\mathcal{E}_{t,i}$ by assigning each of the following with equal probability mass $m_{t,i}^{-1}$.

$$\hat{z}_{t_i(j),i}^t := \min \left\{ 1, \underbrace{v_{t_i(j),i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{t_i(j),i})}_{\text{Optimistic debiased noise sample}} + \underbrace{\text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i})}_{\text{Optimistic estimated mean}} \right\}, \quad (7)$$

where for any context $x \in \mathcal{X}$, any failure probability $\delta \in (0, 1)$, and any $\theta \in \mathbb{R}^d$

$$\text{UCB}_i^t(\theta, x) := x^\top \theta + \beta_i^t(x, \delta), \quad \text{and} \quad \text{LCB}_i^t(\theta, x) := x^\top \theta - \beta_i^t(x, \delta). \quad (8)$$

The construction of $\mathcal{E}_{t,i}$ is *value-optimistic* because $\hat{z}_{t_i(j),i}^t \geq z_{t_i(j),i}^t$ holds. If one has $m_{t,i}$ samples at round t for box i , then $\mathcal{E}_{t,i}$ stochastically dominates the empirical distribution constructed by using samples $\{z_{t_i(j),i}^t\}_{j=1}^{m_{t,i}}$. Then, for each box i , we construct the optimistic distribution $\hat{\mathcal{E}}_{t,i}$ by moving a small amount probability mass for the value-optimistic empirical distribution $\mathcal{E}_{t,i}$ so that it dominates the true distribution (the detailed construction of $\hat{\mathcal{E}}_{t,i}$ and the proof are deferred to Appendix D).

Lemma 4.1. *With probability at least $1 - \delta$, for all $t \in [T]$, $\hat{\mathcal{E}}_t \succeq_{SD} D_t$.*

Once the optimism is ensured, we prove the following regret bound for Algorithm 5 in Appendix D.

Theorem 4.2. *For contextual linear Pandora's Box problem, choosing $\delta = T^{-1}$ for Algorithm 5 ensures $\text{Reg}_T = \tilde{\mathcal{O}}(nd\sqrt{T})$ regret bound.*

While the algorithm of [Atsidakou et al., 2024] applies to our setting with $\tilde{\mathcal{O}}(nT^{5/6})$ regret, Theorem 4.2 gives a $\sqrt{T}$ -type regret bound in the contextual linear setting, which is better whenever $d = o(T^{1/3})$. The tightness of $\tilde{\mathcal{O}}(nd\sqrt{T})$ remains an open question. In the contextual linear bandit setting with $\theta_1 = \dots = \theta_n$, the known lower bound is $\Omega(d\sqrt{T})$ [Dani et al., 2008]. Our setting however requires learning a separate noise distribution for each box, and we conjecture that this makes the linear dependence on n unavoidable, even when all boxes share the same θ .

Analysis. As Lemma 4.1 ensures that $\hat{\mathcal{E}}_t \succeq_{SD} D_t$ for all t , we follow the same argument in the non-contextual case to arrive at Eq. (4). The decomposition of the main term in Eq. (4) slightly differs from that of the non-contextual case. We introduce an empirical distribution $\hat{D}_{t,i}$ for the decomposition, which is constructed by assigning equal probability mass $m_{t,i}^{-1}$ for each $\{z_{t_i(j),i}^t\}_{j=1}^{m_{t,i}}$. Notice that conditioning on history before t , $\hat{D}_{t,i}$ is a product distribution. Then, conditioning on the history before t , for any i , the main term is rewritten as

$$\begin{aligned} \text{Term}_{t,i} &= \underbrace{\mathbb{E}_{z \sim \hat{\mathcal{E}}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim \mathcal{E}_{t,i}} [\tilde{R}_i(\sigma_t; z)]}_{(I)} + \underbrace{\mathbb{E}_{z \sim \mathcal{E}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim \hat{D}_{t,i}} [\tilde{R}_i(\sigma_t; z)]}_{(II)} \\ &\quad + \underbrace{\mathbb{E}_{z \sim \hat{D}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim D_{t,i}} [\tilde{R}_i(\sigma_t; z)]}_{(III)}. \end{aligned}$$

Here, term (I) captures the error of optimistic reweighting, which arises from shifting probability mass in the empirical distribution $\mathcal{E}_{t,i}$ to construct the optimistic distribution $\hat{\mathcal{E}}_{t,i}$. Term (III) quantifies the estimation error between the empirical distribution $\hat{D}_{t,i}$ and the true underlying distribution D_i . Both can be bounded in a similar way as in the non-contextual case. It thus remains to handle (II), which captures the error incurred by the optimistic value shifts. One can show that

$$(II) = \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \left(\tilde{R}(\sigma_t; \hat{z}_{t_i(j),i}) - \tilde{R}(\sigma_t; z_{t_i(j),i}) \right) \leq \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \left| z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t \right|,$$

where the inequality uses 1-Lipschitz property (see [Lemma 3.2.1](#)). Following standard linear bandit analysis (e.g., [[Abbasi-Yadkori et al., 2011](#)]), the total regret incurred by (II) for all boxes is bounded by (omitting constant)

$$\sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[\frac{Q_{t,i}}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \left| z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t \right| \right] \leq \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[Q_{t,i} \left(\beta_i^t(x_{t,i}, \delta) + \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \beta_i^t(x_{t_i(j),i}, \delta) \right) \right],$$

The first summation term can be bounded by

$$\sum_{t=1}^T \sum_{i=1}^n \mathbb{E} [Q_{t,i} \beta_i^t(x_{t,i}, \delta)] \leq \tilde{\mathcal{O}} \left(\sum_{i=1}^n \mathbb{E} \left[\sum_{j=1}^{m_{T,i}} \min\{1, \|x_{t_i(j),i}\|_{V_{t_i(j),i}^{-1}}\} \right] \right) \leq \tilde{\mathcal{O}}(nd\sqrt{T}),$$

where the last inequality applies the Cauchy–Schwarz inequality followed by the elliptical potential lemma [[Abbasi-Yadkori et al., 2011](#), Lemma 11]. The second summation term can be bounded similarly. We refer readers to [Appendix D](#) for the omitted details.

5 Conclusion and Future Work

In this paper, we propose new algorithms for the Pandora’s Box problem in both non-contextual and contextual settings. In the non-contextual case, our algorithm achieves a regret of $\tilde{\mathcal{O}}(\sqrt{nT})$ which matches the known lower bound up to logarithmic factors. For the contextual case, we design a modified algorithm that attains $\tilde{\mathcal{O}}(nd\sqrt{T})$ regret. We further extend our methods and analysis to the Prophet Inequality problem in both settings, achieving similar improvements and regret bounds. Our results also elicit compelling open questions:

1. **Minimax regret for Prophet Inequality.** We establish an upper bound of $\tilde{\mathcal{O}}(\sqrt{nT})$ for Prophet Inequality problem, while [Gatmiry et al. \[2024\]](#) proved a lower bound of $\Omega(\sqrt{T})$, leaving a $\sqrt{n}$ -gap. In fact, [Jin et al. \[2024\]](#) show that the optimal sample complexity of Prophet Inequality is $\tilde{\mathcal{O}}(\frac{1}{\epsilon^2})$, independent of n , which suggests that a simple explore-then-commit strategy achieves an n -independent regret upper bound of $\tilde{\mathcal{O}}(T^{2/3})$. Whether $\tilde{\mathcal{O}}(\sqrt{T})$ is achievable remains open.
2. **Tighter regret for contextual linear case.** In the non-contextual case, our bound scales as $\sqrt{n}$, while in the contextual linear setting, it grows linearly with n . Can this dependence be improved? It also remains unclear whether the linear dependence on the dimension d is improvable.

Acknowledgments and Disclosure of Funding

HL is supported by NSF Award IIS-1943607. LJR and JL are supported in part by ONR YIP N000142012571, and NSF awards AF-2312775, CPS-1844729.

References

- Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, 2011.
- A. Agarwal, R. Ghuge, and V. Nagarajan. Semi-bandit learning for monotone stochastic optimization. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1260–1274. IEEE, 2024.

- S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.
- A. Atsidakou, C. Caramanis, E. Gergatsouli, O. Papadigenopoulos, and C. Tzamos. Contextual pandora’s box. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 10944–10952, 2024.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 2002.
- H. Beyhaghi and L. Cai. Pandora’s problem with nonobligatory inspection: Optimal structure and a ptas. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 803–816, 2023.
- H. Beyhaghi and L. Cai. Recent developments in pandora’s box problem: Variants and applications. *ACM SIGecom Exchanges*, 21(1):20–34, 2024.
- S. Boodaghians, F. Fusco, P. Lazos, and S. Leonardi. Pandora’s box problem with order constraints. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 439–458, 2020.
- S. Chawla, E. Gergatsouli, Y. Teng, C. Tzamos, and R. Zhang. Pandora’s box with correlations: Learning and approximation. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1214–1225. IEEE, 2020.
- V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, 2008.
- T. Ezra, M. Feldman, N. Gravin, and Z. G. Tang. Prophet matching with general arrivals. *Mathematics of Operations Research*, 47(2):878–898, 2022.
- M. Feldman, O. Svensson, and R. Zenklusen. Online contention resolution schemes with applications to bayesian selection problems. *SIAM Journal on Computing*, 50(2):255–300, 2021.
- H. Fu and T. Lin. Learning utilities and equilibria in non-truthful auctions. *Advances in Neural Information Processing Systems*, 33:14231–14242, 2020.
- H. Fu, J. Li, and D. Liu. Pandora box problem with nonobligatory inspection: Hardness and approximation scheme. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 789–802, 2023.
- K. Gattmirey, T. Kesselheim, S. Singla, and Y. Wang. Bandit algorithms for prophet inequality and pandora’s box. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 462–500. SIAM, 2024.
- E. Gergatsouli and C. Tzamos. Online learning for min sum set cover and pandora’s box. In *International Conference on Machine Learning*, pages 7382–7403. PMLR, 2022.
- E. Gergatsouli and C. Tzamos. Weitzman’s rule for pandora’s box with correlations. *Advances in Neural Information Processing Systems*, 36:12644–12664, 2023.
- N. Gravin and H. Wang. Prophet inequality for bipartite matching: Merits of being simple and non adaptive. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 93–109, 2019.
- C. Guo, Z. Huang, and X. Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 662–673, 2019.
- C. Guo, Z. Huang, Z. G. Tang, and X. Zhang. Generalizing complex hypotheses on product distributions: Auctions, prophet inequalities, and pandora’s problem. In *Conference on Learning Theory*, pages 2248–2288. PMLR, 2021.

- J. Jiang, W. Ma, and J. Zhang. Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2022.
- B. Jin, T. Kesselheim, W. Ma, and S. Singla. Sample complexity of posted pricing for a single item. *arXiv preprint arXiv:2406.00819*, 2024.
- R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.
- R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior*, 113:97–115, 2019.
- U. Krengel and L. Sucheston. Semiamarts and finite values. 1977.
- B. Kveton, C. Szepesvari, Z. Wen, and A. Ashkan. Cascading bandits: Learning to rank in the cascade model. In *International conference on machine learning*, pages 767–776. PMLR, 2015.
- T. Lattimore and C. Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, pages 1213–1216, 1984.
- S. Singla. The price of information in combinatorial optimization. In *Proceedings of the twenty-ninth annual ACM-SIAM symposium on discrete algorithms*, pages 2523–2532. SIAM, 2018.
- A. Slivkins. Introduction to multi-armed bandits. *Found. Trends Mach. Learn.*, 12(1-2):1–286, 2019.
- M. Weitzman. *Optimal search for the best alternative*, volume 78. Department of Energy, 1978.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We provide clear motivations for the research questions in the introduction section and explicitly itemize our contributions, outlining how we address and resolve these challenges throughout the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: In this paper, for the contextual model, we only study the linear structure of rewards for both Pandora's Box and Prophet Inequality problems.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All assumptions are clearly stated in [Section 2](#) and we defer the proof to Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: This paper is theoretically oriented. There are no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: This paper is theoretically oriented. There are no experiments.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: This paper is theoretically oriented. There are no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: This paper is theoretically oriented. There are no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: This paper is theoretically oriented. There are no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have reviewed NeurIPS Code of Ethics and commit to it.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In the introduction, we highlight that our research can benefit sequential decision-making problems, modeled by Pandora's Box and Prophet Inequality. We do not foresee any negative societal impacts arising from this work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper is theoretically oriented and does not release any data or models.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: This is a theory-oriented paper with no experiments or code.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This is a theory-oriented paper. There are no new assets introduced.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This is a theory-oriented paper. There are no experiments or research with human objects in this paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This is a theory-oriented paper. There are no experiments or research with human objects in this paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: This work does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Appendix

A	Problem Setting of Prophet Inequality	21
B	Related Work	21
C	Omitted Details from the Non-Contextual Settings	23
C.1	Technical Lemmas	23
C.2	Lipschitz Property, Monotonicity, and Sharp Derivative for Pandora’s Box	24
C.3	Lipschitz Property, Monotonicity, and Sharp Derivative for Prophet Inequality	27
C.4	Proof of $\tilde{O}(\sqrt{nT})$ Regret Bound for Pandora’s Box	28
C.5	$\tilde{O}(\sqrt{nT})$ Regret Bound for Prophet Inequality	31
C.6	High-Probability Regret Bounds	32
C.7	Extension to Subgaussian Rewards	32
D	Omitted Details of Contextual Linear Setting	33
D.1	Preliminaries	33
D.2	Algorithms and $\tilde{O}(nd\sqrt{T})$ Bounds for Pandora’s Box and Prophet Inequality	34
D.3	Proof of Lemma D.2.2	36

Algorithm 3 Generic threshold algorithm for Prophet Inequality

Input: threshold $\sigma = (\sigma_1, \dots, \sigma_n)$.

The environment generates an unknown realization $v = (v_1, \dots, v_n)$.

for $i = 1, \dots, n$ **do**

 Open box i to observe v_i .

 If $v_i \geq \sigma_i$ or $i = n$, then stop and take reward v_i .

A Problem Setting of Prophet Inequality

Online Prophet Inequality. In this problem, the learner is given n boxes, each with an unknown fixed reward distribution D_i supported on $[0, 1]$. The learner plays a T -round repeated game where at each round $t \in [T]$, the learner inspects boxes in a fixed order. Upon opening a box i , the learner receives a reward $v_{t,i} \in [0, 1]$ independently sampled from D_i . Then, she needs to decide immediately whether to accept this reward and stop the process at round t or discard this reward and open the next box. The utility for the round is the chosen reward.

Contextual Linear Prophet Inequality. This setting extends the non-contextual one, and the assumptions on reward distributions and boundedness are the same as [Section 2](#).

Optimal Policy and Regret. To measure the performance, we compare our algorithm against the per-round optimal policy. Let $D_t = (D_{t,1}, \dots, D_{t,n})$ be the product distribution of boxes. The optimal policy for Prophet Inequality in round t is also a *threshold* policy, in which the threshold for each box i , denoted by $\sigma_{t,i}^*$ is derived from a *reverse/backward programming*.

Definition A.1 (Optimal Threshold on D_t). Let $\sigma_{t,n}^* = \mathbb{E}_{X \sim D_{t,n}}[X]$. For $i = n-1, \dots, 1$, let $\sigma_{t,i}^* = \mathbb{E}_{X \sim D_{t,i}}[\max\{X, \sigma_{t,i+1}^*\}]$.

The optimal policy at round t is to run [Algorithm 3](#) with threshold vector σ_t^* . In fact, the optimal threshold of box i computed from [Definition A.1](#) is equal to the expected reward obtained by running $\sigma_{t,i}^*, \dots, \sigma_{t,n}^*$ directly from box i [[Gatmiry et al., 2024](#)].

More formally, similar to those definitions in Pandora’s Box setting, for any product distribution $D = (D_1, \dots, D_n)$, let $\sigma_D = (\sigma_{D,1}, \dots, \sigma_{D,n})$ denote the optimal threshold vector, where each $\sigma_{D,i}$ is the threshold calculated by reverse programming. For any threshold vector $\sigma = (\sigma_1, \dots, \sigma_n) \in [0, 1]^n$ and realization $v = (v_1, \dots, v_n) \in [0, 1]^n$, we use $S(\sigma; v) \in [n]$ corresponds to the box where [Algorithm 3](#) stops. Further, we define the utility function and expected utility function, respectively:

$$R(\sigma; v) = v_{S(\sigma; v)}, \quad \text{and} \quad R(\sigma; D) = \mathbb{E}_{v \sim D}[R(\sigma; v)]. \quad (9)$$

With these definitions, the optimal expected utility is $R(\sigma_{D_t}; D_t)$ since when $v \sim D_t$, $S(\sigma_{D_t}; v) \in [n]$ corresponds to the box chosen by the optimal policy. Let $\sigma_t \in \mathbb{R}^n$ denote the threshold vector selected by the algorithm at round t . Throughout the paper, all our algorithms in the Prophet Inequality setting follow the general structure outlined in [Algorithm 3](#) using σ_t as input at round t , and thus the expected utility of our algorithms at round t is $R(\sigma_t; D_t)$. The cumulative regret in T rounds Reg_T is then defined as

$$\text{Reg}_T = \mathbb{E} \left[\sum_{t=1}^T R(\sigma_{D_t}; D_t) - R(\sigma_t; D_t) \right],$$

where $D_t = D$ for the non-contextual case.

B Related Work

The Pandora’s Box problem, introduced by [Weitzman \[1978\]](#), has been extended in various directions, including matroid constraints [[Singla, 2018](#)], correlated distributions [[Chawla et al., 2020](#), [Gergatsoulis and Tzamos, 2023](#)], order constraints [[Boodaghians et al., 2020](#)], and settings without inspection [[Fu et al., 2023](#), [Beyhaghi and Cai, 2023](#)]. The Prophet Inequality problem was first studied by [Krengel and Sucheston \[1977\]](#), [Samuel-Cahn \[1984\]](#), with subsequent generalizations to k -selection [[Alaei, 2014](#), [Jiang et al., 2022](#)], matroid constraints [[Kleinberg and Weinberg, 2012, 2019](#), [Feldman et al., 2021](#)], and online matching [[Gravin and Wang, 2019](#), [Ezra et al., 2022](#)]. These works typically assume known distributions and focus on achieving constant-factor approximation guarantees.

A growing body of work has begun to study online variants of the Prophet Inequality and Pandora's Box problems [Guo et al., 2019, Fu and Lin, 2020, Guo et al., 2021, Garmiry et al., 2024, Agarwal et al., 2024]. In the PAC setting, Fu and Lin [2020], Guo et al. [2021] establish a sample complexity upper bound of $\tilde{O}(\frac{n}{\epsilon^2})$ for Pandora's Box, matching the known lower bound up to logarithmic factors, where ϵ denotes the desired accuracy. Guo et al. [2021] further develop a general framework for learning over product distributions, which applies to both problems and yields a sample complexity of $\tilde{O}(\frac{n}{\epsilon^2})$ for Prophet Inequality as well. Later, Jin et al. [2024] provide a near-optimal sample complexity $\tilde{O}(\frac{1}{\epsilon^2})$ for Prophet Inequality, which removes the linear dependence on n . For the regret minimization problem, Garmiry et al. [2024] establish $\tilde{O}(n^3\sqrt{T})$ regret bound for Prophet Inequality and $\tilde{O}(n^{4.5}\sqrt{T})$ regret bound for Pandora's Box under a constrained feedback model, where the learner observes only the final reward without knowing which box produced it. Subsequently, Agarwal et al. [2024] introduce a general framework for monotone stochastic optimization under semi-bandit feedback, achieving $\tilde{O}(n\sqrt{T})$ regret for both problems. All these works assume a fixed but unknown product distribution. Moreover, despite this progress, the best known regret bounds for both Prophet Inequality and Pandora's Box under semi-bandit feedback remain at $\tilde{O}(n\sqrt{T})$, which falls short of the known lower bounds [Garmiry et al., 2024].

To better capture the dynamics in decision-making scenarios, Gergatsouli and Tzamos [2022] study the Pandora's Box problem in the adversarial setting where the reward of each box is chosen by an adversary and obtain sublinear regret bounds by comparing the algorithm against a weak benchmark. Later, Garmiry et al. [2024] show that in the same setting, when the algorithm competes against the optimal policy, no algorithm can achieve sublinear regret, even with full information feedback. To make the problem tractable while still allowing distributions to change over time, Atsidakou et al. [2024] explore a contextual Pandora's Box model where in each round, the learner observes a context vector for each box, whose optimal threshold can be approximated by a function of the context and an unknown vector. They provide a reduction from contextual Pandora's Box to an instance of online regression. Their algorithm suffers a $\tilde{O}(T^{5/6})$ regret bound if the threshold can be linearly approximated. We note that their algorithm for the linear case also applies to our setting (see the discussion below), but our algorithm enjoys $\tilde{O}(nd\sqrt{T})$, which improves on the aforementioned bound whenever $d = o(T^{1/3})$.

Discussion of [Atsidakou et al., 2024]. Here, we show how the algorithm proposed by Atsidakou et al. [2024] can be applied to our setting. Their work makes the following realizability assumption.

Assumption 1 (Realizability of [Atsidakou et al., 2024]). *There exists $w_1, \dots, w_n \in \mathbb{R}^d$ and a function h such that every time $t \in [T]$ and every box $i \in [n]$, the optimal threshold for distribution $D_{t,i}$ is equal to $h(w_i, x'_{t,i})$, i.e., $\sigma_{t,i}^* = h(w_i, x'_{t,i})$ where $x'_{t,i}$ is a context vector of box i at round t and*

$$\mathbb{E}_{X \sim D_{t,i}} [(X - \sigma_{t,i}^*)_+] = c_i.$$

Below, we show that in our setting, the optimal threshold can be realized by a linear function h (that is, $h(a, b) = a^\top b$). To see this, first note that the optimal threshold $\sigma_{t,i}^*$ in our setting satisfies $\mathbb{E}_{X \sim D_{t,i}} [(X - \sigma_{t,i}^*)_+] = c_i$, which, based on the definition of $D_{t,i}$, is equivalent to $\mathbb{E}_{X \sim D_i} [(X - (\sigma_{t,i}^* - \mu_{t,i}))_+] = c_i$ where $\mu_{t,i} = \theta_i^\top x_{t,i}$. Therefore, if we let σ_i^* be the unique solution of $\mathbb{E}_{X \sim D_i} [(X - \sigma_i^*)_+] = c_i$, then we have

$$\sigma_{t,i}^* = \mu_{t,i} + \sigma_i^* = \langle w_i, x'_{t,i} \rangle = h(w_i, x'_{t,i}),$$

for $w_i = (\theta_{i,1}, \dots, \theta_{i,d}, \sigma_1, \dots, \sigma_n) \in \mathbb{R}^{d+n}$, and $x'_{t,i} = ([x_{t,i}]_1, \dots, [x_{t,i}]_d, 0, \dots, 1, \dots, 0) \in \mathbb{R}^{d+n}$, where $x'_{t,i}$ takes value one at the $d+i$ index. This proves that the optimal threshold is realized by a linear function of some context, making the algorithm of [Atsidakou et al., 2024] applicable to our setting.

A seemingly related strand is cascading bandits [Kveton et al., 2015], where the learner presents an ordered list of items to a user, with items corresponding to boxes in our setting. Then, the user scans the list and clicks the first attractive item if any. The learner observes binary click signals for items up to that click, which yields order dependent semi bandit feedback. Despite this shared feedback structure, Pandora's Box and Prophet Inequality differ in who controls how many items are inspected. In cascading bandits the number is user driven since observation stops at the first click. In Pandora's

Box and in Prophet Inequality the number is policy driven since the algorithm adaptively opens boxes and decides when to stop.

C Omitted Details from the Non-Contextual Settings

C.1 Technical Lemmas

Lemma C.1.1 (Bernstein inequality). *Given $m \in \mathbb{N}$, let $X_1, X_2, \dots, X_m$ be i.i.d. random variables such that $\mathbb{E}[X_i] = 0$, $\mathbb{E}[X_i^2] = \sigma^2$, and $|X_i| \leq M$ for some constant $M > 0$. Then, for all $t > 0$*

$$\mathbb{P}\left(\left|\sum_{i=1}^m X_i\right| > t\right) \leq 2 \exp\left(-\frac{t^2}{2m\sigma^2 + \frac{2}{3}Mt}\right).$$

Lemma C.1.2. *Let $D = (D_1, \dots, D_n)$ be a $[0, 1]$ -bounded product distribution, and for each i , let $\mathcal{E}_{N_i, i}$ be the empirical distribution constructed by assigning N_i i.i.d. samples from D_i with equal probability mass $1/N_i$. With probability at least $1 - \delta$, for any $i \in [n]$, any $N_i \in [T]$, and any $z \in [0, 1]$,*

$$|F_{D_i}(z) - F_{\mathcal{E}_{N_i, i}}(z)| \leq \sqrt{\frac{2F_{D_i}(z)(1 - F_{D_i}(z)) \log(2T^2 n \delta^{-1})}{N_i}} + \frac{\log(2T^2 n \delta^{-1})}{N_i}. \quad (10)$$

Proof. This proof follows a similar idea to Guo et al. [2019, Lemma 5], but we apply an additional union bound to account for all possible N_i . Fix i and N_i . It suffices to show that for all $z \in [0, 1]$ such that $F_{D_i}(z)$ is multiples of $1/N_i$, the following holds.

$$|F_{D_i}(z) - F_{\mathcal{E}_{N_i, i}}(z)| \leq \sqrt{\frac{2F_{D_i}(z)(1 - F_{D_i}(z)) \log(2T^2 n \delta^{-1})}{N_i}} + \frac{2 \log(2T^2 n \delta^{-1})}{3N_i}, \quad (11)$$

This is because given the results above, the general bound on the difference of CDF of any value differs at most an extra additive factor $\frac{1}{N_i} \leq \frac{\log(2T^2 n \delta^{-1})}{3N_i}$.

Then fix any z such that $F_{D_i}(z)$ is multiples of $1/N_i$. Let Y_i be the indicator function of the i -th sample no smaller than z and let $X_i = Y_i - \mathbb{E}[Y_i]$. By Bernstein inequality (see Lemma C.1.1), taking t be N_i times the RHS of Eq. (11),

$$\mathbb{P}\left(\left|\sum_{i=1}^{N_i} X_i\right| > t\right) = \exp\left(-\frac{t^2}{2N_i F_{D_i}(z)(1 - F_{D_i}(z)) + \frac{2}{3}t}\right) \leq \frac{\delta}{nT^2}$$

Since $N_i \leq T$, by a union bound over all N_i , with probability at least $1 - \frac{\delta}{nT}$, Eq. (11) holds for all z such that $F_{D_i}(z)$ is multiples of $1/N_i$. Further, by union bounds over all $i \in [n]$ and $N_i \in [T]$, we get that Eq. (10) does not hold with probability at most δ . $\square$

As $F_{D_i}(\cdot)$ is unknown, we present the following corollary, which provides a confidence bound depending on known $F_{\mathcal{E}_i}(\cdot)$.

Corollary C.1.3. *Under the same setting of Lemma C.1.2, suppose that the concentration bounds of Lemma C.1.2 all hold, and then for any $i \in [n]$, any $N_i \in [T]$ and any $x \in [0, 1]$,*

$$|F_{D_i}(x) - F_{\mathcal{E}_{N_i, i}}(x)| \leq \sqrt{\frac{2F_{\mathcal{E}_{N_i, i}}(x)(1 - F_{\mathcal{E}_{N_i, i}}(x))L}{N_i}} + \frac{L}{N_i}, \text{ where } L = 4 \log(2nT^2 \delta^{-1}). \quad (12)$$

Proof. Fix any i and N_i . For notional simplicity, we write $\mathcal{E}_i = \mathcal{E}_{N_i, i}$. For any $x \in [0, 1]$,

$$\begin{aligned} F_{D_i}(x)(1 - F_{D_i}(x)) &\leq F_{\mathcal{E}_i}(x)(1 - F_{\mathcal{E}_i}(x)) + |F_{D_i}(x)(1 - F_{D_i}(x)) - F_{\mathcal{E}_i}(x)(1 - F_{\mathcal{E}_i}(x))| \\ &\leq F_{\mathcal{E}_i}(x)(1 - F_{\mathcal{E}_i}(x)) + |F_{D_i}(x) - F_{\mathcal{E}_i}(x)|, \end{aligned} \quad (13)$$

where the last inequality uses the fact that function $y - y^2$ is 1-Lipschitz for $y \in [0, 1]$.

Let $L' = \log(2nT^2\delta^{-1})$. From [Eq. \(10\)](#), we can show for any $x \in [0, 1]$

$$\begin{aligned}
|F_{D_i}(x) - F_{\mathcal{E}_i}(x)| &\leq \sqrt{\frac{2F_{D_i}(x)(1 - F_{D_i}(x))L'}{N_i}} + \frac{L'}{N_i} \\
&\leq \sqrt{\frac{2L'(F_{\mathcal{E}_i}(x)(1 - F_{\mathcal{E}_i}(x)) + |F_{D_i}(x) - F_{\mathcal{E}_i}(x)|)}{N_i}} + \frac{L'}{N_i} \\
&\leq \sqrt{\frac{2L'F_{\mathcal{E}_i}(x)(1 - F_{\mathcal{E}_i}(x))}{N_i}} + \sqrt{\frac{2L'|F_{D_i}(x) - F_{\mathcal{E}_i}(x)|}{N_i}} + \frac{L'}{N_i} \\
&\leq \sqrt{\frac{2L'F_{\mathcal{E}_i}(x)(1 - F_{\mathcal{E}_i}(x))}{N_i}} + \frac{|F_{D_i}(x) - F_{\mathcal{E}_i}(x)|}{2} + \frac{2L'}{N_i}, \tag{14}
\end{aligned}$$

where the second inequality uses [Eq. \(13\)](#), and the last inequality uses $2\sqrt{ab} \leq a + b$ for any $a, b \geq 0$ (here $a = L'/N_i$ and $b = \frac{1}{2}|F_{D_i}(x) - F_{\mathcal{E}_i}(x)|$). Rearranging the above inequality gives the claimed bound for fixed i, N_i . Since this argument holds for all i, N_i , the lemma thus follows. $\square$

Lemma C.1.4. *For all t, i , $F_{\hat{\mathcal{E}}_{t,i}}(\cdot)$ constructed by [Eq. \(1\)](#) is a valid CDF.*

Proof. To show $F_{\hat{\mathcal{E}}_{t,i}}(\cdot)$ is valid, we show that it is a right-continuous, non-decreasing function in the range of $[0, 1]$. Consider any given t, i . From [Eq. \(1\)](#), one can easily see $F_{\hat{\mathcal{E}}_{t,i}}(x) \in [0, 1]$ for all $x \in [0, 1]$. As $F_{\mathcal{E}_{t,i}}(\cdot)$ is right-continuous and max function preserves the right-continuous property, $F_{\hat{\mathcal{E}}_{t,i}}(\cdot)$ is again right-continuous.

Now, we verify the monotonicity of $F_{\hat{\mathcal{E}}_{t,i}}(\cdot)$. Since $F_{\mathcal{E}_{t,i}}(\cdot)$ is non-decreasing, we only need to show $F_{\hat{\mathcal{E}}_{t,i}}(x)$ is non-decreasing as a function of $F_{\mathcal{E}_{t,i}}(x)$. Let $k = \frac{L}{m_{t,i}}$. For the interval of x such that $F_{\mathcal{E}_{t,i}}(x) < k$, $F_{\hat{\mathcal{E}}_{t,i}}(x)$ is 0, and thus it is non-decreasing. Then we only need to consider the interval of x such that $F_{\mathcal{E}_{t,i}}(x) \geq k$. We let $y = F_{\mathcal{E}_{t,i}}(x)$, and rewrite the formulation for $F_{\hat{\mathcal{E}}_{t,i}}(x)$ as $f(y) = y - \sqrt{2ky(1-y)} - k$. In this case, $f'(y) = 1 - \frac{\sqrt{k(1-2y)}}{\sqrt{2y(1-y)}}$. When $y \geq \frac{1}{2}$, this is obviously non-negative, and thus $f(y)$ is non-decreasing. Then, we consider the case of $y < \frac{1}{2}$ and show that $\sqrt{2y(1-y)} > \sqrt{k(1-2y)}$. As $y \geq k$, it is enough to show $\sqrt{2y(1-y)} > \sqrt{y(1-2y)}$. This is equivalent to show $4y^2 - 2y - 1 < 0$ when $y < \frac{1}{2}$, which can be easily verified.

As the argument holds for all t, i , the lemma thus follows. $\square$

Lemma C.1.5 (Restatement of [Lemma 3.1.2](#)). *With probability at least $1 - \delta$, for all $t \in [T]$, $\hat{\mathcal{E}}_t \succeq_{SD} D$.*

Proof. For this proof, we show $\hat{\mathcal{E}}_{t,i} \succeq_{SD} D_i$ for all t, i . Fix t, i . It suffices to show that for all $x \in [0, 1]$, $F_{\hat{\mathcal{E}}_{t,i}}(x) \leq F_{D_i}(x)$. For $x = 1$, $F_{\hat{\mathcal{E}}_{t,i}}(x) = F_{D_i}(x) = 1$. For $x \in [0, 1)$, we have

$$\begin{aligned}
F_{\hat{\mathcal{E}}_{t,i}}(x) &= \max \left\{ 0, F_{\mathcal{E}_{t,i}}(x) - \sqrt{\frac{2F_{\mathcal{E}_{t,i}}(x)(1 - F_{\mathcal{E}_{t,i}}(x))L}{m_{t,i}}} - \frac{L}{m_{t,i}} \right\} \\
&\leq \max \{0, F_{D_i}(x)\} = F_{D_i}(x),
\end{aligned}$$

where the inequality follows from [Corollary C.1.3](#).

Repeating the same argument for all $i \in [n]$ completes the proof. $\square$

C.2 Lipschitz Property, Monotonicity, and Sharp Derivative for Pandora's Box

Before proving Lipschitz property and monotonicity for $\tilde{R}_i(\sigma_t, z)$, we first introduce the following supporting results.

Definition C.2.1 ((n, D, c) -Pandora's box instance). We use (n, D, c) to characterize a Pandora's box instance where product distribution $D = (D_1, \dots, D_n)$ is over n boxes, each of which is associated with a cost c_i .

For any (n, D, c) -Pandora's box instance, let $W_u : [0, 1]^n \times [0, 1]^n \rightarrow 2^{[n]}$ be a general threshold rule with initial value u . Specifically, for any initial value u , threshold σ , and realization $x \sim D$, $W_u(\sigma, x)$ is the set of opened boxes following rule W_u . For any u, σ , if the initial value $u \geq \max_i \sigma_i$, then $W_u(\sigma, \cdot) = \emptyset$, that is, the learner does not open any box and keeps the initial value u . Otherwise, the learner keeps the initial value u in hand but continues to open boxes by running [Algorithm 1](#) with $V_{\max} = u$ and threshold σ . Let

$$R(\sigma; x; u; W_v) = \max \left\{ u, \max_{i \in W_v(\sigma; x)} x_i \right\} - \sum_{i \in W_v(\sigma; x)} c_i$$

$$R(\sigma; D; u; W_v) = \mathbb{E}_{x \sim D} [R(\sigma; x; u; W_v)].$$

Here, $R(\sigma; x; u; W_v)$ is the utility if the learner adopts rule W_v to run threshold σ with initial value u over realization x . As shown by [Weitzman \[1978\]](#), for any initial value u , the expected utility $R(\sigma_D; D; u; W_u)$ enjoys the optimality.

Lemma C.2.2 (Monotonicity and 1-Lipschitz Property for Pandora's Box). For any initial values u, v , if the learner runs general threshold rules W_u and W_v with $u \geq v$, respectively on a (n, D, c) -Pandora's Box instance, then we have

$$0 \leq \mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_u)] - \mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_v)] \leq u - v.$$

Consequently, the following 1-Lipschitz property holds: for any u, v ,

$$\left| \mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_u)] - \mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_v)] \right| \leq |u - v|.$$

Proof. Consider any $u \geq v$. For any realization $x = (x_1, \dots, x_n)$ drawn from product distribution D , the utility given the initial value s ($= u$ or v) and rule W is

$$R(\sigma_D; x; s; W) = \max \left\{ s, \max_{i \in W(\sigma_D; x)} x_i \right\} - \sum_{i \in W(\sigma_D; x)} c_i.$$

Taking $W = W_v$, and $s = u$ and v , respectively, we have

$$\begin{aligned} & \mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_v)] - \mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_v)] \\ &= \mathbb{E}_{x \sim D} \left[\max \left\{ u, \max_{i \in W_u(\sigma_D; x)} x_i \right\} - \max \left\{ v, \max_{i \in W_v(\sigma_D; x)} x_i \right\} \right] \\ &\geq 0. \end{aligned}$$

Given the initial value u and threshold σ_D , using rule W_u recovers the Weitzman's optimal algorithm. Thus, the corresponding expected utility should be no smaller than that of running rule W_v with the same threshold σ_D and initial value u , i.e., $\mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_v)] \leq \mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_u)]$, which implies $0 \leq \mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_u)] - \mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_v)]$.

Taking $W = W_u$, and $s = u$ and v , respectively, we have

$$\begin{aligned} & \mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_u)] - \mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_u)] \\ &= \mathbb{E}_{x \sim D} \left[\max \left\{ u, \max_{i \in W_u(\sigma_D; x)} x_i \right\} - \max \left\{ v, \max_{i \in W_u(\sigma_D; x)} x_i \right\} \right] \\ &\leq u - v. \end{aligned}$$

where the inequality holds since for any x , $f_x(v) = \max\{v, x\}$ is 1-Lipschitz with respect to v .

As W_v is optimal given initial value v and threshold σ_D , we have $\mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_u)] \leq \mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_v)]$, implying $\mathbb{E}_{x \sim D} [R(\sigma_D; x; u; W_u)] - \mathbb{E}_{x \sim D} [R(\sigma_D; x; v; W_v)] \leq u - v$. Thus, the lemma follows. $\square$

For any $x \in [0, 1]^n$, we define $x_{>i} := (x_{i+1}, \dots, x_n)$ and $x_{<i} := (x_1, \dots, x_{i-1})$. Based on the definition of $\mathcal{H}_{t,i} = (D_1, \dots, D_i, \hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n})$, we define

$$\mathcal{H}_{t,>i} := (\hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n}) \quad \text{and} \quad \mathcal{H}_{t,<i} := (D_1, \dots, D_{i-1}).$$

For any $z \in [0, 1]$, any $x \in [0, 1]^n$, and any $\sigma \in [0, 1]^n$, we define

$$\begin{aligned} \tilde{R}_i(\sigma; x_{<i}, z) &:= \mathbb{E}_{x_{>i} \sim \mathcal{H}_{t,>i}} [R(\sigma; (x_1, \dots, x_{i-1}, z, x_{i+1}, \dots, x_n))] \\ &= \mathbb{E}_{x_{>i} \sim (\hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n})} [R(\sigma; (x_1, \dots, x_{i-1}, z, x_{i+1}, \dots, x_n))], \end{aligned} \quad (15)$$

where the expectation is taken only for $x_{>i} := (x_{i+1}, \dots, x_n)$.

Lemma C.2.3 (Restatement of [Lemma 3.2.1](#)). *For all $t \in [T]$ and $i \in [n]$, the map $\tilde{R}_i(\sigma_t; z)$ is 1-Lipschitz and monotonically-increasing with respect to z in the Pandora's Box problem.*

Proof. For simplicity, we assume without loss of generality that $\sigma_{t,1} \geq \sigma_{t,2} \geq \dots \geq \sigma_{t,n}$. For any $z \in [0, 1]$ and any $x \in [0, 1]^n$, if box i is opened at round t , then the map $\tilde{R}_i(\sigma_t; x_{<i}, z)$ is the expected utility of running σ_t on $(x_{i+1}, \dots, x_n) \sim \mathcal{H}_{t,>i}$ with initial value $\max\{x_1, \dots, x_{i-1}, z\}$. A key fact here is that $\sigma_t = \sigma_{\hat{\mathcal{E}}_t}$ and $\mathcal{H}_{t,>i} = (\hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n})$, which imply that the algorithm runs the optimal threshold in descending order. Thus [Lemma C.2.2](#) implies that $\tilde{R}_i(\sigma_t; x_{<i}, z)$ is 1-Lipschitz with respect to z and for any $u \geq v$, and that $\tilde{R}_i(\sigma_t; x_{<i}, u) - \tilde{R}_i(\sigma_t; x_{<i}, v) \in [0, u - v]$. This, in turn, implies that

$$\tilde{R}_i(\sigma_t; u) - \tilde{R}_i(\sigma_t; v) = \mathbb{E}_{x_{<i} \sim (D_1, \dots, D_{i-1})} [\tilde{R}_i(\sigma_t; x_{<i}, u) - \tilde{R}_i(\sigma_t; x_{<i}, v) \mid E_{\sigma_t, i}] \in [0, u - v].$$

The 1-Lipschitz property of $\tilde{R}_i(\sigma; z)$ is an immediate corollary, and the proof is thus complete. $\square$

Lemma C.2.4 (Restatement of [Lemma 3.2.2](#)). *Suppose that $\sigma_{t,1} \geq \sigma_{t,2} \geq \dots \geq \sigma_{t,n}$. The Pandora's Box problem satisfies the following:*

$$0 \leq \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) \leq \frac{\prod_{j<i} F_{D_j}(z)}{Q_{t,i}}, \quad \forall z \in [0, \sigma_{t,i+1}).$$

Proof. According to [Lemma 3.2.1](#), the map $\tilde{R}_i(\sigma_t; z)$ is monotonically increasing with respect to z , which implies that $0 \leq \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z)$. Then, we only need to prove for any $0 \leq v \leq u < \sigma_{t,i+1}$,

$$\tilde{R}_i(\sigma_t; u) - \tilde{R}_i(\sigma_t; v) \leq (u - v) \cdot \frac{\prod_{j<i} F_{D_j}(u)}{Q_{t,i}}. \quad (16)$$

To this end, we deduce the following:

$$\begin{aligned} &\tilde{R}_i(\sigma_t; u) - \tilde{R}_i(\sigma_t; v) \\ &= \mathbb{E}_{x_{<i} \sim (D_1, \dots, D_{i-1})} [\tilde{R}_i(\sigma_t; x_{<i}, u) - \tilde{R}_i(\sigma_t; x_{<i}, v) \mid E_{\sigma_t, i}] \\ &\stackrel{(a)}{=} \mathbb{E}_{x_{<i} \sim (D_1, \dots, D_{i-1})} [\tilde{R}_i(\sigma_t; x_{<i}, \max\{x_{<i}, u\}) - \tilde{R}_i(\sigma_t; x_{<i}, \max\{x_{<i}, v\}) \mid E_{\sigma_t, i}] \\ &\leq \mathbb{E}_{x_{<i} \sim (D_1, \dots, D_{i-1})} [\max\{x_{<i}, u\} - \max\{x_{<i}, v\} \mid E_{\sigma_t, i}] \\ &\stackrel{(b)}{=} \mathbb{E}_{x_{<i} \sim (D_1, \dots, D_{i-1})} \left[(\max\{x_{<i}, u\} - \max\{x_{<i}, v\}) \cdot \mathbb{I}\left\{u > \max_{j<i} x_j\right\} \mid E_{\sigma_t, i} \right] \\ &\leq (u - v) \mathbb{E}_{x_{<i} \sim (D_1, \dots, D_{i-1})} \left[\mathbb{I}\left\{u > \max_{j<i} x_j\right\} \mid E_{\sigma_t, i} \right] \\ &= (u - v) \cdot \mathbb{P}_{x_{<i} \sim (D_1, \dots, D_{i-1})} \left(\max_{j<i} x_j < u \mid E_{\sigma_t, i} \right) \end{aligned}$$

$$\begin{aligned}
&= (u - v) \cdot \frac{\mathbb{P}_{x_{<i} \sim (D_1, \dots, D_{i-1})}(\max_{j < i} x_j < u, E_{\sigma_t, i})}{\mathbb{P}_{x_{<i} \sim (D_1, \dots, D_{i-1})}(E_{\sigma_t, i})} \\
&\stackrel{(c)}{=} (u - v) \cdot \frac{\mathbb{P}_{x_{<i} \sim (D_1, \dots, D_{i-1})}(\max_{j < i} x_j < u)}{\mathbb{P}_{x_{<i} \sim (D_1, \dots, D_{i-1})}(E_{\sigma_t, i})} \\
&= (u - v) \frac{\prod_{j < i} F_{D_j}(u)}{Q_{t, i}},
\end{aligned}$$

where (a) holds since conditioning on $E_{\sigma_t, i}$ (i.e., opening box i), the expected reward remains unchanged by replacing u (or v) by the maximum so far, the inequality uses the 1-Lipschitz property of $\tilde{R}_i(\sigma_t; x_{<i}, z)$ with respect to variable z , (b) follows from the fact that $\max(x_{<i}, u) - \max(x_{<i}, v)$ is non-zero only when $\max_{j < i} x_j < u$ as $u \geq v$, and (c) holds due to $\{\max_{j < i} x_j < u\} \subseteq E_{\sigma_t, i}$ where this containment follows from the fact that $u < \sigma_{t, i+1} \leq \sigma_{t, i}$ and $E_{\sigma_t, i} = \{\max_{j < i} x_j < \sigma_{t, i}\}$.

Thus, the stated claim holds. $\square$

C.3 Lipschitz Property, Monotonicity, and Sharp Derivative for Prophet Inequality

In the following, we prove the similar results for Prophet Inequality.

Lemma C.3.1. *Consider any two product distributions $D_u = (u, D_2, D_3, \dots, D_n)$ and $D_v = (v, D_2, D_3, \dots, D_n)$, where the first box deterministically produces rewards u and v , respectively, and the remaining boxes share the same distributions. If $u \geq v$, then*

$$0 \leq R(\sigma_{D_u}; D_u) - R(\sigma_{D_v}; D_v) \leq u - v.$$

As a corollary, we have the following 1-Lipschitz property. For any u, v ,

$$|R(\sigma_{D_u}; D_u) - R(\sigma_{D_v}; D_v)| \leq |u - v|.$$

Proof. As σ_{D_u} is computed by backward induction from [Definition A.1](#), we have $R(\sigma_{D_u}; D_u) = \sigma_{D_u, 1} = \max\{u, \sigma_{D_u, 2}\}$. Moreover, since the distribution of boxes from $2, \dots, n$ are the same, $\sigma_{D_u, i} = \sigma_{D_v, i}$ for all $i \geq 2$. For any $u \geq v$,

$$R(\sigma_{D_u}; D_u) - R(\sigma_{D_v}; D_v) = \max\{u, \sigma_{D_u, 2}\} - \max\{v, \sigma_{D_v, 2}\} \in [0, u - v],$$

where the inequality holds since max function is 1-Lipschitz and $\sigma_{D_u, 2} = \sigma_{D_v, 2}$. Thus, the lemma follows. $\square$

In the Prophet Inequality problem, the definition of $\tilde{R}_i(\sigma; z)$ takes the same form as in [Eq. \(6\)](#), but uses the utility function R defined in [Eq. \(9\)](#). Similarly, we follow [Eq. \(15\)](#) to define $\tilde{R}_i(\sigma; x_{<i}, z)$, again using the utility function defined in [Eq. \(9\)](#).

Lemma C.3.2. *For all $t \in [T]$, $i \in [n]$, $\tilde{R}_i(\sigma; z)$ is 1-Lipschitz and monotonically-increasing with respect to z in Prophet Inequality problem.*

Proof. As the order of the boxes is fixed across time, we use [Lemma C.3.1](#) to repeat a similar argument in [Lemma 3.2.1](#) to complete the proof. The only difference is that conditioning on event that the learner opens box i , $\tilde{R}_i(\sigma; x_{<i}, z)$ can be interpreted as the expected utility of running σ_t on $(x_{i+1}, \dots, x_n) \sim \mathcal{H}_{t, > i}$ with $x_i = z$ since $x_1, \dots, x_{i-1}$ are discarded. $\square$

Lemma C.3.3 (Sharp Bound of Derivative for Prophet Inequality). *The Prophet Inequality problem satisfies*

$$\frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) = 0, \quad \forall z \in [0, \sigma_{t, i}].$$

Proof. For any $0 \leq v \leq u < \sigma_{t, i}$, we have that

$$\tilde{R}_i(\sigma_t; u) - \tilde{R}_i(\sigma_t; v) = \mathbb{E}_{x_{<i} \sim (D_1, \dots, D_{i-1})} \left[\tilde{R}_i(\sigma_t; x_{<i}, u) - \tilde{R}_i(\sigma_t; x_{<i}, v) \mid E_{\sigma_t, i} \right] = 0,$$

where the second equality uses the fact that if the learner opens box i and the realized value is smaller than $\sigma_{t, i}$, then they will discard it and keep opening subsequent boxes. Thus, the lemma follows. $\square$

C.4 Proof of $\tilde{\mathcal{O}}(\sqrt{nT})$ Regret Bound for Pandora's Box

The analysis in this section conditions on the event that the concentration bounds of [Lemma C.1.2](#) hold with respect to D and $\mathcal{E}_t$ for all t . Recall from [Section 3.2](#) that regret is bounded by

$$\text{Reg}_T \leq \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[Q_{t,i} \cdot \underbrace{\left(\mathbb{E}_{z \sim \hat{\mathcal{E}}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim D_i} [\tilde{R}_i(\sigma_t; z)] \right)}_{=: \text{Term}_{t,i}} \right]. \quad (17)$$

Consider any given round t and suppose, without loss of generality, that

$$\sigma_{t,1} \geq \sigma_{t,2} \geq \dots \geq \sigma_{t,n}.$$

Since $\tilde{R}_i(\sigma_t; z)$ is 1-Lipschitz with respect to z , the map $\tilde{R}_i(\sigma_t; z)$ is absolutely continuous on $[0, 1]$, which implies that it is differentiable almost everywhere in $[0, 1]$. By the fundamental theorem of calculus, for any $x \in [0, 1]$, we have that $\tilde{R}_i(\sigma_t; x) = \tilde{R}_i(\sigma_t; 0) + \int_0^x \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz$. Therefore, we have that

$$\begin{aligned} \text{Term}_{t,i} &= \mathbb{E}_{z \sim \hat{\mathcal{E}}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim D_i} [\tilde{R}_i(\sigma_t; z)] \\ &= \mathbb{E}_{z \sim \hat{\mathcal{E}}_{t,i}} \left[\tilde{R}_i(\sigma_t; 0) + \int_0^z \frac{\partial}{\partial t} \tilde{R}_i(\sigma_t; t) dt \right] - \mathbb{E}_{z \sim D_i} \left[\tilde{R}_i(\sigma_t; 0) + \int_0^z \frac{\partial}{\partial t} \tilde{R}_i(\sigma_t; t) dt \right] \\ &= \mathbb{E}_{z \sim \hat{\mathcal{E}}_{t,i}} \left[\int_0^1 \frac{\partial}{\partial t} \tilde{R}_i(\sigma_t; t) \mathbb{I}\{z \geq t\} dt \right] - \mathbb{E}_{z \sim D_i} \left[\int_0^1 \frac{\partial}{\partial t} \tilde{R}_i(\sigma_t; t) \mathbb{I}\{z \geq t\} dt \right] \\ &= \int_0^1 (1 - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz - \int_0^1 (1 - F_{D_i}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz. \end{aligned}$$

According to the analysis in [Section 3.2](#), we bound

$$\text{Term}_{t,i} \leq \underbrace{\int_{\sigma_{t,i+1}}^1 |F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| dz}_{A_{t,i}} + \underbrace{\int_0^{\sigma_{t,i+1}} (F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz}_{B_{t,i}}.$$

Now, we complete the bounds on $A_{t,i}$ and $B_{t,i}$.

Bounding $A_{t,i}$. We have

$$\begin{aligned} A_{t,i} &\leq \int_{\sigma_{t,i+1}}^1 |F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| dz + \int_{\sigma_{t,i+1}}^1 |F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| dz \\ &\leq \mathcal{O} \left(\frac{L}{m_{t,i}} + \int_{\sigma_{t,i+1}}^1 \sqrt{\frac{F_{D_i}(z)(1 - F_{D_i}(z))L}{m_{t,i}}} dz + \int_{\sigma_{t,i+1}}^1 \sqrt{\frac{F_{\mathcal{E}_{t,i}}(z)(1 - F_{\mathcal{E}_{t,i}}(z))L}{m_{t,i}}} dz \right) \\ &\leq \mathcal{O} \left(\frac{L}{m_{t,i}} + \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} + \sqrt{\frac{(1 - F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1}))L}{m_{t,i}}} \right) \\ &\leq \mathcal{O} \left(\sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right), \end{aligned} \quad (18)$$

where the second inequality bounds $|F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)|$ by [Lemma C.1.2](#) and bounds $|F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)|$ by the construction [Eq. \(1\)](#), and the third inequality bounds $1 - F_{D_i}(z) \leq 1 - F_{D_i}(\sigma_{t,i+1})$ (similar for $1 - F_{\mathcal{E}_{t,i}}(z)$) as $z \geq \sigma_{t,i+1}$. The last inequality follows from the following reasoning:

$$\begin{aligned} &\sqrt{\frac{(1 - F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1}))L}{m_{t,i}}} \\ &\leq \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L + L |F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1}) - F_{D_i}(\sigma_{t,i+1})|}{m_{t,i}}} \end{aligned}$$

$$\begin{aligned}
&\leq \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} + \sqrt{\frac{L |F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1}) - F_{D_i}(\sigma_{t,i+1})|}{m_{t,i}}} \\
&\leq \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} + \frac{|F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1}) - F_{D_i}(\sigma_{t,i+1})|}{2} + \frac{L}{m_{t,i}} \\
&\leq \mathcal{O} \left(\sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right), \tag{19}
\end{aligned}$$

where the third inequality uses $\sqrt{2ab} \leq a + b$ for any $a, b \geq 0$, and the last inequality bounds $|F_{\mathcal{E}_{t,i}}(\sigma_{t,i+1}) - F_{D_i}(\sigma_{t,i+1})|$ by [Lemma C.1.2](#).

Bounding $B_{t,i}$. By [Lemma 3.2.2](#), we bound

$$\begin{aligned}
B_{t,i} &\leq \int_0^{\sigma_{t,i+1}} |F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| \frac{\sqrt{\prod_{j<i} F_{D_j}(z)} \sqrt{\prod_{j<i} F_{D_j}(z)}}{Q_{t,i}} dz \\
&\leq \int_0^{\sigma_{t,i+1}} \left(|F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)| + |F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| \right) \sqrt{\frac{\prod_{j<i} F_{D_j}(z)}{Q_{t,i}}} dz, \tag{20}
\end{aligned}$$

where the second inequality bounds $\prod_{j<i} F_{D_j}(z) \leq \prod_{j<i} F_{D_j}(\sigma_{t,i+1}) \leq \prod_{j<i} F_{D_j}(\sigma_{t,i}) = Q_{t,i}$ for all $z < \sigma_{t,i+1}$.

On the one hand, we use [Lemma C.1.2](#) to bound for any z

$$|F_{D_i}(z) - F_{\mathcal{E}_{t,i}}(z)| \leq \mathcal{O} \left(\sqrt{\frac{F_{D_i}(z)(1 - F_{D_i}(z))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right) \leq \mathcal{O} \left(\sqrt{\frac{(1 - F_{D_i}(z))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right).$$

On the other hand, we use the construction of optimistic distribution $\hat{\mathcal{E}}_{t,i}$ in [Eq. \(1\)](#) to show for any z

$$|F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| \leq \mathcal{O} \left(\sqrt{\frac{(1 - F_{\mathcal{E}_{t,i}}(z))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right) \leq \mathcal{O} \left(\sqrt{\frac{(1 - F_{D_i}(z))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right),$$

where the last inequality repeats the same argument in [Eq. \(19\)](#).

Plugging the two bounds above into [Eq. \(20\)](#), we have that

$$\begin{aligned}
B_{t,i} &\leq \mathcal{O} \left(\int_0^{\sigma_{t,i+1}} \left(\sqrt{\frac{(1 - F_{D_i}(z))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right) \sqrt{\frac{\prod_{j<i} F_{D_j}(z)}{Q_{t,i}}} dz \right) \\
&= \mathcal{O} \left(\int_0^{\sigma_{t,i+1}} \left(\sqrt{\frac{\prod_{j<i} F_{D_j}(z)(1 - F_{D_i}(z))L}{Q_{t,i}m_{t,i}}} + \sqrt{\frac{\prod_{j<i} F_{D_j}(z)}{Q_{t,i}}} \frac{L}{m_{t,i}} \right) dz \right) \\
&\leq \mathcal{O} \left(\frac{L}{m_{t,i}} + \int_0^{\sigma_{t,i+1}} \left(\sqrt{\frac{\prod_{j<i} F_{D_j}(z)(1 - F_{D_i}(z))L}{Q_{t,i}m_{t,i}}} \right) dz \right) \\
&\leq \mathcal{O} \left(\frac{L}{m_{t,i}} + \int_0^1 \left(\sqrt{\frac{\prod_{j<i} F_{D_j}(z)(1 - F_{D_i}(z))L}{Q_{t,i}m_{t,i}}} \right) dz \right),
\end{aligned}$$

where the second inequality follows from

$$\sqrt{\frac{\prod_{j<i} F_{D_j}(z)}{Q_{t,i}}} \leq \sqrt{\frac{\prod_{j<i} F_{D_j}(\sigma_{t,i+1})}{Q_{t,i}}} \leq \sqrt{\frac{\prod_{j<i} F_{D_j}(\sigma_{t,i})}{Q_{t,i}}} = 1$$

based on the fact that $z \leq \sigma_{t,i+1}$ and $Q_{t,i} = \prod_{j<i} F_{D_j}(\sigma_{t,i})$.

Regret for round t . From the bounds of $A_{t,i}$ and $B_{t,i}$, the regret at round t is bounded as

$$\begin{aligned} \text{Reg}(t) &\leq \sum_{i \in [n]} \text{Term}_{t,i} \leq \mathcal{O} \left(\mathbb{E} \left[\sum_{i \in [n]} Q_{t,i} \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} \right. \right. \\ &\quad \left. \left. + \sum_{i \in [n]} \int_0^1 \sqrt{\frac{\prod_{j < i} F_{D_j}(z)(1 - F_{D_i}(z))LQ_{t,i}}{m_{t,i}}} dz + \frac{LQ_{t,i}}{m_{t,i}} \right] \right). \end{aligned}$$

The first term is bounded as follows:

$$\begin{aligned} &\mathbb{E} \left[\sum_{i \in [n]} Q_{t,i} \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i+1}))L}{m_{t,i}}} \right] \\ &\leq \mathbb{E} \left[\sqrt{\sum_{i \in [n]} Q_{t,i}(1 - F_{D_i}(\sigma_{t,i+1}))} \sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}} \right] \leq \mathbb{E} \left[\sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}} \right] =, \end{aligned}$$

where the first inequality uses the Cauchy–Schwarz inequality, and the second inequality follows from the the fact that

$$\begin{aligned} \sqrt{\sum_{i \in [n]} Q_{t,i}(1 - F_{D_i}(\sigma_{t,i+1}))} &= \sqrt{\sum_{i \in [n]} \prod_{j < i} F_{D_j}(\sigma_{t,i})(1 - F_{D_i}(\sigma_{t,i+1}))} \\ &\leq \sqrt{\sum_{i \in [n]} \prod_{j < i} F_{D_j}(\sigma_{t,j+1})(1 - F_{D_i}(\sigma_{t,i+1}))} \\ &= \sqrt{\sum_{i \in [n]} \left(\prod_{j < i} F_{D_j}(\sigma_{t,j+1}) - \prod_{j \leq i} F_{D_j}(\sigma_{t,j+1}) \right)} \\ &\leq 1. \end{aligned}$$

Here, the first inequality holds since the descending sorted assumption on σ_t gives $\sigma_{t,j+1} \geq \sigma_{t,i}$ for all $j < i$, and the last inequality follows from the fact that the summation forms a telescoping sum.

To bound the second term, we show that

$$\begin{aligned} &\sum_{i \in [n]} \int_0^1 \sqrt{\frac{\prod_{j < i} F_{D_j}(z)(1 - F_{D_i}(z))LQ_{t,i}}{m_{t,i}}} dz \\ &= \int_0^1 \sum_{i \in [n]} \sqrt{\prod_{j < i} F_{D_j}(z)(1 - F_{D_i}(z))} \sqrt{\frac{LQ_{t,i}}{m_{t,i}}} dz \\ &\leq \int_0^1 \sqrt{\sum_{i \in [n]} \prod_{j < i} F_{D_j}(z)(1 - F_{D_i}(z))} \sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}} dz \\ &\leq \sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}}, \end{aligned}$$

where the first inequality uses the Cauchy–Schwarz inequality, and the last inequality follows from the fact that $\sum_{i \in [n]} \prod_{j < i} F_{D_j}(z)(1 - F_{D_i}(z))$ forms a telescoping sum.

Combining all bounds above, we have

$$\text{Reg}(t) \leq \mathcal{O} \left(\mathbb{E} \left[\sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}} + \frac{LQ_{t,i}}{m_{t,i}} \right] \right).$$

Algorithm 4 Proposed algorithm for Prophet Inequality

Input: confidence $\delta \in (0, 1)$, horizon T .

Initialize: open all boxes once and observe rewards. Set $m_{1,i} = 1$ for all $i \in [n]$.

for $t = 1, 2, \dots, T$ **do**

 For each $i \in [n]$, construct an optimistic distribution $\hat{\mathcal{E}}_{t,i}$ according to [Eq. \(1\)](#).

 Use [Definition A.1](#) with distribution $\hat{\mathcal{E}}_t$ to compute $\sigma_t = (\sigma_{t,1}, \dots, \sigma_{t,n})$.

 Run [Algorithm 3](#) with σ_t to open a set of boxes, denoted by $\mathcal{B}_t$ and observe rewards $\{v_{t,i}\}_{i \in \mathcal{B}_t}$.

 Update counters $m_{t+1,i} = m_{t,i} + 1, \forall i \in \mathcal{B}_t$ and $m_{t+1,i} = m_{t,i}, \forall i \notin \mathcal{B}_t$.

Summing $\text{Reg}(t)$ over all t . The Cauchy–Schwarz inequality gives

$$\begin{aligned} \text{Reg}_T &\leq \mathcal{O} \left(\mathbb{E} \left[\sum_{t=1}^T \left(\sqrt{\sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}}} + \frac{LQ_{t,i}}{m_{t,i}} \right) \right] \right) \\ &\leq \mathcal{O} \left(\mathbb{E} \left[\sqrt{T \sum_{i \in [n]} \sum_{t=1}^T \frac{LQ_{t,i}}{m_{t,i}}} + \sum_{t=1}^T \frac{LQ_{t,i}}{m_{t,i}} \right] \right) \\ &\leq \mathcal{O} \left(\sqrt{T \sum_{i \in [n]} \mathbb{E} \left[\sum_{t=1}^T \frac{LQ_{t,i}}{m_{t,i}} \right]} + \mathbb{E} \left[\sum_{t=1}^T \frac{LQ_{t,i}}{m_{t,i}} \right] \right), \end{aligned} \quad (21)$$

where the last inequality uses Jensen’s inequality.

Finally, it remains to bound $\mathbb{E} \left[\sum_{t=1}^T LQ_{t,i}/m_{t,i} \right]$. Let $I_{t,i} = \mathbb{I}\{E_{\sigma_t,i}\}$.

$$\mathbb{E} \left[\sum_{t=1}^T \sum_{i \in [n]} \frac{LQ_{t,i}}{m_{t,i}} \right] = \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in [n]} \frac{L\mathbb{E}_t[I_{t,i}]}{m_{t,i}} \right] = \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in [n]} \frac{LI_{t,i}}{m_{t,i}} \right] \leq \mathcal{O}(nL \log(T)), \quad (22)$$

where the second equality holds due to $m_{t,i}$ is deterministic given history before round t .

Combining the bound $\mathbb{E} \left[\sum_{t=1}^T LQ_{t,i}/m_{t,i} \right] = \mathcal{O}(nL \log T)$ with [Eq. \(21\)](#) completes the proof of [Theorem 3.1.3](#).

C.5 $\tilde{\mathcal{O}}(\sqrt{nT})$ Regret Bound for Prophet Inequality

The proposed algorithm for the Prophet Inequality problem is shown in [Algorithm 4](#).

Theorem C.5.1. *For the Prophet Inequality problem, choosing $\delta = T^{-1}$ for [Algorithm 4](#) ensures*

$$\text{Reg}_T = \mathcal{O} \left(\sqrt{nT \log(T) \log(nT)} + n \log(T) \log(nT) \right).$$

Proof. The analysis conditions on the event that the concentration bounds of [Lemma C.1.2](#) hold with respect to D and $\mathcal{E}_t$ for all t . As shown by [Guo et al. \[2021\]](#), the Prophet Inequality satisfies the same monotonicity as Pandora’s Box. Thus, the analysis follows the same approach as used in the Pandora’s Box problem to bound $\text{Reg}_T \leq \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} [Q_{t,i} \cdot \text{Term}_{t,i}]$ where $\text{Term}_{t,i}$ has the same definition as in [Eq. \(4\)](#) with a new utility function $\tilde{R}$ defined in [Eq. \(9\)](#).

As shown in [Lemma C.3.2](#), the Prophet Inequality problem has both monotonicity and smoothness properties. Therefore, repeating the same analysis in [Section 3.2](#) (splitting the integral on $\sigma_{t,i}$ instead of $\sigma_{t,i+1}$) gives

$$\begin{aligned} \text{Term}_{t,i} &\leq \int_{\sigma_{t,i}}^1 \left| F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z) \right| dz + \int_0^{\sigma_{t,i}} \left(F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z) \right) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz \\ &= \underbrace{\int_{\sigma_{t,i}}^1 \left| F_{D_i}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z) \right| dz}_{A_{t,i}}, \end{aligned}$$

where the equality applies [Lemma C.3.3](#). We use the same argument as in [Eq. \(18\)](#) to bound $A_{t,i}$: indeed,

$$A_{t,i} \leq \mathcal{O} \left(\sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i}))L}{m_{t,i}}} + \frac{L}{m_{t,i}} \right).$$

Thus, the regret at round t is bounded by

$$\text{Reg}(t) \leq \mathcal{O} \left(\sum_{i \in [n]} Q_{t,i} \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i}))L}{m_{t,i}}} + \sum_{i \in [n]} \frac{Q_{t,i}L}{m_{t,i}} \right).$$

In particular, we bound

$$\begin{aligned} & \sum_{i \in [n]} Q_{t,i} \sqrt{\frac{(1 - F_{D_i}(\sigma_{t,i}))L}{m_{t,i}}} \\ & \leq \sqrt{\sum_{i \in [n]} Q_{t,i}(1 - F_{D_i}(\sigma_{t,i}))} \sqrt{\sum_{i \in [n]} \frac{Q_{t,i}L}{m_{t,i}}} \\ & = \sqrt{\sum_{i \in [n]} \prod_{j < i} F_{D_j}(\sigma_{t,j})(1 - F_{D_i}(\sigma_{t,i}))} \sqrt{\sum_{i \in [n]} \frac{Q_{t,i}L}{m_{t,i}}} \leq \sqrt{\sum_{i \in [n]} \frac{Q_{t,i}L}{m_{t,i}}}, \end{aligned}$$

where the equality holds since $Q_{t,i} = \prod_{j < i} F_{D_j}(\sigma_{t,j})$ in the Prophet Inequality setting.

By repeating the same argument as in [Eq. \(21\)](#) and [Eq. \(22\)](#), the proof of [Theorem C.5.1](#) is complete. $\square$

C.6 High-Probability Regret Bounds

The regret metric now is

$$\text{Reg}'_T := \sum_{t=1}^T (R(\sigma_D; D) - R(\sigma_t; D)). \quad (23)$$

Following the same argument in [Section 3.2](#), we arrive at [Eq. \(4\)](#) without expectation. As σ_t and $H_{t,i}$ are deterministic given history, we have

$$R(\sigma_t; H_{t,i-1}) - R(\sigma_t; H_{t,i}) = \mathbb{E}_t[R(\sigma_t; H_{t,i-1}) - R(\sigma_t; H_{t,i})],$$

where $\mathbb{E}_t[\cdot]$ is the conditional expectation given history before round t . Then, we can follow the same analysis to arrive at $\sqrt{\sum_{i,t} \frac{Q_{t,i}}{m_{t,i}}}$. To handle this, let $M_{t,i} = \frac{Q_{t,i} - I_{t,i}}{m_{t,i}}$, and $\{M_{t,i}\}_t$ is martingale difference sequence. By Freedman's inequality, with probability at least $1 - \delta$, we have $\sum_t M_{t,i} \leq \sqrt{2V \log(1/\delta)} + \log(1/\delta)$ where $V = \sum_{t=1}^T \mathbb{E}_t[M_{t,i}^2]$. Notice that $V = \frac{Q_{t,i} - Q_{t,i}^2}{m_{t,i}} \leq \frac{Q_{t,i}}{m_{t,i}}$. Thus, we have $\sum_t M_{t,i} \leq \sqrt{2 \log(1/\delta) \sum_t Q_{t,i}/m_{t,i}} + \log(1/\delta) \leq \frac{1}{2} \sum_t Q_{t,i}/m_{t,i} + 2 \log(1/\delta)$ where the last step uses the AM-GM inequality. Rearranging it gives $\sum_{t=1}^T \frac{Q_{t,i}}{m_{t,i}} \leq O(\log(1/\delta) + \sum_{t=1}^T \frac{I_{t,i}}{m_{t,i}})$. Therefore, by choosing δ properly and applying a union bound, with high probability, the regret is bounded by $\tilde{O}(\sqrt{nT})$.

C.7 Extension to Subgaussian Rewards

For each box, if the reward distribution is sub-Gaussian, then our algorithms still work with minor modifications. Specifically, if the reward distribution of box i is K_i -sub-Gaussian, one can set a range $[-K_i \sqrt{2 \log(2T^2 n)}, K_i \sqrt{2 \log(2T^2 n)}]$. Then, the algorithm updates parameters only if the received reward of each box i is within range of $[-K_i \sqrt{2 \log(2T^2 n)}, K_i \sqrt{2 \log(2T^2 n)}]$. On those rounds, the algorithm and its regret analysis coincide exactly with the bounded-reward setting. Meanwhile, by a union bound on sub-Gaussian tails, the probability that any reward ever exceeds its range is at most $1/T$. Therefore, its contribution to the expected regret is only $\tilde{O}(1)$.

D Omitted Details of Contextual Linear Setting

D.1 Preliminaries

Lemma D.1.1 (DKW Inequality). *Given a natural number N , let $X_1, \dots, X_N$ be i.i.d. samples from distribution D with cumulative distribution function $F_D(\cdot)$. Let $F_{\mathcal{E}}(\cdot)$ be the associated empirical distribution function $\mathcal{E}$ such that for all $x \in \mathbb{R}$, we define $F_{\mathcal{E}}(x) := \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{X_i \leq x\}$. For every $\varepsilon > 0$,*

$$\mathbb{P} \left(\sup_{x \in \mathbb{R}} |F_{\mathcal{E}}(x) - F_D(x)| > \varepsilon \right) \leq 2 \exp(-2N\varepsilon^2).$$

Lemma D.1.2. *Let $D = (D_1, \dots, D_n)$ be a product distribution, and for each i , let $\mathcal{E}_{N_i, i}$ be the empirical distribution built from N_i i.i.d. samples of D_i . With probability at least $1 - \delta/2$, for any $i \in [n]$, any $N_i \in [T]$,*

$$\sup_{z \in \mathbb{R}} |F_{D_i}(z) - F_{\mathcal{E}_{N_i, i}}(z)| \leq \sqrt{\frac{\log(4Tn\delta^{-1})}{2N_i}}. \quad (24)$$

Proof. We first fix any i and N_i and apply [Lemma D.1.1](#). Then, the lemma follows by using a union bound over all $i \in [n]$ and $N_i \in [T]$. $\square$

Lemma D.1.3. *With probability at least $1 - \delta/2$, for any $t > 0$ and any $i \in [n]$,*

$$\left\| \theta_i - \hat{\theta}_{t,i} \right\|_{V_{t,i}} \leq \alpha_\delta, \quad \text{where} \quad \alpha_\delta = 1 + \sqrt{2 \log(2n/\delta) + d \log \left(1 + \frac{T}{d} \right)}.$$

Proof. Fix any $i \in [n]$. Since the noise distribution of each box is $\frac{1}{4}$ -subgaussian, and $\{x_{t,i}\}_t$ is a $\mathbb{R}^d$ -valued stochastic process such that $x_{t,i}$ is measurable with respect to the history, [[Lattimore and Szepesvári, 2020](#), Theorem 20.5] directly gives the claimed result for the fixed i . Using a union bound over all $i \in [n]$ completes the proof. $\square$

Lemma D.1.4 (Elliptical Potential Lemma) (see, e.g., Lemma 11 in [[Abbasi-Yadkori et al., 2011](#)])). *Let $I \in \mathbb{R}^{d \times d}$ be an identity matrix and a sequence of vectors $a_1, \dots, a_N \in \mathbb{R}^d$ with $\|a_i\|_2 \leq 1$ for all $i \in [N]$. Let $V_t = I + \sum_{s \leq t} a_s a_s^\top$. Then, for all $N \in \mathbb{N}$, we have*

$$\sum_{t=1}^N \min \left\{ 1, \|a_t\|_{V_{t-1}^{-1}}^2 \right\} \leq 2d \log \left(1 + \frac{N}{d} \right).$$

Definition D.1.5 (Definition of $\tilde{D}_t$). *Let $\tilde{D}_t = (\tilde{D}_{t,1}, \dots, \tilde{D}_{t,n})$ where each $\tilde{D}_{t,i}$ is an empirical distribution which assigns $m_{t,i}$ i.i.d. noise samples from D_i with equal probability mass $1/m_{t,i}$.*

Definition D.1.6 ("Nice" event). *Let $\mathcal{E}$ be the nice event that the concentration bounds in [Lemma D.1.2](#) for $D, \tilde{D}_t$ hold for all $t \in [T]$, and the concentration bounds in [Lemma D.1.3](#) hold simultaneously.*

Lemma D.1.7. *Suppose that $\mathcal{E}$ holds. For any context $x \in \mathcal{X}$ and any round $t \in [T]$,*

$$\left| \left\langle x, \hat{\theta}_{t,i} - \theta_i \right\rangle \right| \leq \beta_i^{t+1}(x).$$

Proof. For any t, i, x , one can show that

$$\left| \left\langle x, \hat{\theta}_{t,i} - \theta_i \right\rangle \right| \leq \|x\|_{V_{t,i}^{-1}} \left\| \theta_i - \hat{\theta}_{t,i} \right\|_{V_{t,i}} \leq \alpha_\delta \|x\|_{V_{t,i}^{-1}} = \beta_i^{t+1}(x),$$

where the first inequality follows from the Cauchy–Schwarz inequality and the second inequality uses [Lemma D.1.3](#). $\square$

D.2 Algorithms and $\tilde{O}(nd\sqrt{T})$ Bounds for Pandora's Box and Prophet Inequality

The omitted algorithm for both the Pandora's Box and Prophet Inequality problems is presented in [Algorithm 5](#). Unlike the non-contextual case (see [Eq. \(1\)](#)), we use a slightly different construction for the optimistic distribution. This is because the Bernstein-type construction requires access to the CDF of an empirical distribution based on i.i.d. samples, which is unavailable in our setting. As a result, we use the following approach instead.

Optimistic distribution $\hat{\mathcal{E}}_{t,i}$. Let $L = \frac{1}{2} \log(4nT/\delta)$. The CDF of $\hat{\mathcal{E}}_{t,i}$ is constructed as follows:

$$F_{\hat{\mathcal{E}}_{t,i}}(x) = \begin{cases} 1, & x = 1; \\ \max \left\{ 0, F_{\mathcal{E}_{t,i}}(x) - \sqrt{\frac{L}{m_{t,i}}} \right\}, & 0 \leq x < 1. \end{cases} \quad (25)$$

Now, we prove [Lemma 4.1](#) which shows that the optimistic distribution $\hat{\mathcal{E}}_t$ stochastically dominates D_t for all t .

Lemma D.2.1 (Restatement of [Lemma 4.1](#)). *Suppose that $\mathcal{E}$ holds. For all $t \in [T]$, $\hat{\mathcal{E}}_t \succeq_{SD} D_t$.*

Proof. For this proof, we show $\hat{\mathcal{E}}_{t,i} \succeq_{SD} D_{t,i}$ for all i, t . Fix a box $i \in [n]$ and a round $t \in [T]$. If $c \geq 1$, then

$$\mathbb{P}_{X \sim \hat{\mathcal{E}}_{t,i}}(X \leq c) = \mathbb{P}_{X \sim D_{t,i}}(X \leq c) = 1.$$

For any $c < 1$, we have that

$$\begin{aligned} \mathbb{P}_{X \sim \hat{\mathcal{E}}_{t,i}}(X \leq c) &= \max \left\{ 0, \mathbb{P}_{X \sim \mathcal{E}_{t,i}}(X \leq c) - \sqrt{\frac{L}{m_{t,i}}} \right\} \\ &\leq \max \left\{ 0, \mathbb{P}_{Z \sim D_{t,i}}(Z \leq c) \right\} \\ &= \mathbb{P}_{Z \sim D_{t,i}}(Z \leq c), \end{aligned}$$

The inequality holds due to the following reason:

$$\begin{aligned} &\mathbb{P}_{X \sim \mathcal{E}_{t,i}}(X \leq c) - \sqrt{\frac{L}{m_{t,i}}} \\ &= \frac{1}{m} \sum_{j=1}^m \mathbb{I} \left\{ \eta_{t_i(j),i} + \mu_{t_i(j),i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{t_i(j),i}) + \text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) \leq c \right\} - \sqrt{\frac{L}{m_{t,i}}} \\ &\leq \frac{1}{m} \sum_{j=1}^m \mathbb{I} \left\{ \eta_{t_i(j),i} + \text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) \leq c \right\} - \sqrt{\frac{L}{m_{t,i}}} \\ &\leq \frac{1}{m_{t,i}} \sum_{j=1}^m \mathbb{I} \left\{ \eta_{t_i(j),i} \leq c - \mu_{t,i} \right\} - \sqrt{\frac{L}{m_{t,i}}} \\ &\leq \mathbb{P}_{X \sim D_i}(X \leq c - \mu_{t,i}) \\ &= \mathbb{P}_{Z \sim D_{t,i}}(Z \leq c), \end{aligned}$$

where the first inequality follows from the facts that $\mu_{t_i(j),i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{t_i(j),i}) \geq 0$, the second inequality uses the fact that $\text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) \geq \mu_{t,i}$, and the last inequality follows from the concentration bound given by $\mathcal{E}$.

Conditioning on the "nice" event, the argument holds for all i and t . The proof is complete. $\square$

Next, we present the following lemma, whose proof is in the next subsection.

Lemma D.2.2. *Suppose that $\mathcal{E}$ holds. For both the Pandora's Box and Prophet Inequality problems, we have that*

$$\text{Reg}_T \leq \sum_{t=1}^T \sum_{i=1}^n \mathcal{O} \left(\mathbb{E} \left[Q_{t,i} \left(\sqrt{\frac{L}{m_{t,i}}} + \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} |z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t| \right) \right] \right).$$

Algorithm 5 Proposed algorithm for contextual linear model

Input: confidence $\delta \in (0, 1)$, horizon T .

Initialize: observe context $\{x_{0,i}\}_{i \in [n]}$ for each box, open each box, observe rewards $\{v_{0,i}\}_{i \in [n]}$, and update counters $\{m_{1,i}\}_{i \in [n]}$ and parameters $\{\hat{\theta}_{0,i}\}_{i \in [n]}$ where $\hat{\theta}_{0,i} = (I + x_{0,i}x_{0,i}^\top)^{-1}x_{0,i}v_{0,i}$.

for $t = 1, 2, \dots, T$ **do**

 Observe contexts $(x_{t,1}, \dots, x_{t,n})$ and compute $\hat{\mu}_{t,i} = \langle \hat{\theta}_{t-1,i}, x_{t,i} \rangle$ for each $i \in [n]$.

 For each $i \in [n]$, use Eq. (25) to construct an empirical distribution $\hat{\mathcal{E}}_{t,i}$ by using samples

$$\left\{ v_{s,i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{s,i}) + \text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) \right\}_{s \leq t: i \in \mathcal{B}_s}.$$

if *Pandora's Box* **then**

 Compute $\sigma_t = (\sigma_{t,1}, \dots, \sigma_{t,n})$ where $\sigma_{t,i}$ is the solution of $\mathbb{E}_{x \sim \hat{\mathcal{E}}_{t,i}}[(x - \sigma_{t,i})_+] = c_i$.

 Run Algorithm 1 with σ_t to open a set of boxes, denoted by $\mathcal{B}_t$ to observe rewards $\{v_{t,i}\}_{i \in \mathcal{B}_t}$.

else if *Prophet Inequality* **then**

 Use Definition A.1 with distribution $\hat{\mathcal{E}}_t$ to compute $\sigma_t = (\sigma_{t,1}, \dots, \sigma_{t,n})$.

 Run Algorithm 3 with σ_t to open a set of boxes, denoted by $\mathcal{B}_t$ to observe rewards $\{v_{t,i}\}_{i \in \mathcal{B}_t}$.

 Update counters $m_{t+1,i} = m_{t,i} + 1, \forall i \in \mathcal{B}_t$ and $m_{t+1,i} = m_{t,i}, \forall i \notin \mathcal{B}_t$.

 Update estimate for each $i \in \mathcal{B}_t$:

$$V_{t,i} = I + \sum_{s \leq t: i \in \mathcal{B}_s} x_{s,i}x_{s,i}^\top \quad \text{where} \quad \hat{\theta}_{t,i} = V_{t,i}^{-1} \sum_{s \leq t: i \in \mathcal{B}_s} x_{s,i}v_{s,i}.$$

With Lemma D.2.2 in hand, we are now ready to prove Theorem 4.2.

Proof of Theorem 4.2. The analysis conditions on event $\mathcal{E}$ which is defined in Definition D.1.6 and occurs with probability at least $1 - \delta$. Recall that $I_{t,i} = \mathbb{I}\{E_{\sigma_t,i}\}$. For the first term, we take the iterated expectation and use the facts that $Q_{t,i} = \mathbb{E}_t[I_{t,i}]$ and $m_{t,i}$ is deterministic given the history to get

$$\sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[Q_{t,i} \sqrt{\frac{L}{m_{t,i}}} \right] = \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[I_{t,i} \sqrt{\frac{L}{m_{t,i}}} \right] \leq \mathcal{O}(n\sqrt{TL}).$$

Next, we bound the summation of the second term. To this end, we show that

$$\begin{aligned} & \left| z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t \right| \\ &= \min \left\{ 1, \left| z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t \right| \right\} \\ &= \min \left\{ 1, \left| \eta_{t_i(j),i} + \mu_{t,i} - \min \left\{ 1, v_{t_i(j),i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{t_i(j),i}) + \text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) \right\} \right| \right\} \\ &\leq \min \left\{ 1, \left| v_{t_i(j),i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{t_i(j),i}) + \text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) - \eta_{t_i(j),i} - \mu_{t,i} \right| \right\} \\ &\leq \min \left\{ 1, \left| v_{t_i(j),i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{t_i(j),i}) - \eta_{t_i(j),i} \right| + \left| \text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) - \mu_{t,i} \right| \right\} \\ &= \min \left\{ 1, \left| \mu_{t_i(j),i} - \text{LCB}_i^t(\hat{\theta}_{t-1,i}, x_{t_i(j),i}) \right| + \left| \text{UCB}_i^t(\hat{\theta}_{t-1,i}, x_{t,i}) - \mu_{t,i} \right| \right\} \\ &= \min \left\{ 1, \left| \langle x_{t_i(j),i}, \theta_i - \hat{\theta}_{t-1,i} \rangle + \beta_i^t(x_{t_i(j),i}) \right| + \left| \langle x_{t,i}, \theta_i - \hat{\theta}_{t-1,i} \rangle + \beta_i^t(x_{t,i}) \right| \right\} \\ &\leq \min \left\{ 1, 2\beta_i^t(x_{t_i(j),i}) + 2\beta_i^t(x_{t,i}) \right\} \\ &\leq \min \left\{ 1, 2\beta_i^t(x_{t_i(j),i}) \right\} + \min \left\{ 1, 2\beta_i^t(x_{t,i}) \right\}, \end{aligned} \tag{26}$$

where the first equality holds due to $z_{t_i(j),i}^t, \hat{z}_{t_i(j),i}^t \in [0, 1]$, the first inequality uses $|b - \min\{1, a\}| \leq |a - b|$ for $b \in [0, 1]$, and the third inequality follows from Lemma D.1.7.

Using [Eq. \(26\)](#), we have that

$$\begin{aligned} & \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[\frac{Q_{t,i}}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \left| z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t \right| \right] \\ & \leq 2 \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[Q_{t,i} \left(\min \{1, \beta_i^t(x_{t,i}, \delta)\} + \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \min \{1, \beta_i^t(x_{t_i(j),i}, \delta)\} \right) \right]. \end{aligned}$$

For the second term on the right hand side, we have that

$$\begin{aligned} & \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[\frac{Q_{t,i}}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \beta_i^t(x_{t_i(j),i}, \delta) \right] \\ & \leq \alpha_\delta \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[\frac{Q_{t,i}}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \min \left\{ 1, \|x_{t_i(j),i}\|_{V_{t-1,i}^{-1}} \right\} \right] \\ & \leq \alpha_\delta \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[\frac{Q_{t,i}}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \min \left\{ 1, \|x_{t_i(j),i}\|_{V_{t_i(j)-1,i}^{-1}} \right\} \right] \\ & \leq \alpha_\delta \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[\frac{Q_{t,i}}{m_{t,i}} \sqrt{m_{t,i} \sum_{j=1}^{m_{t,i}} \min \left\{ 1, \|x_{t_i(j),i}\|_{V_{t_i(j)-1,i}^{-1}}^2 \right\}} \right] \\ & \leq \mathcal{O} \left(\alpha_\delta \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[Q_{t,i} \sqrt{\frac{d \log(1 + T/d)}{m_{t,i}}} \right] \right) \\ & \leq \mathcal{O} \left(\alpha_\delta \cdot n \sqrt{dT \log(1 + T/d)} \right), \end{aligned}$$

where the second inequality follows from $V_{t-1,i} \succeq V_{t_i(j)-1,i}$ for all t and i , the third inequality uses Cauchy–Schwarz inequality, and the fourth inequality follows from [Lemma D.1.4](#).

Analogously, we take the iterated expectation and use the fact that $Q_{t,i} = \mathbb{E}_t[I_{t,i}]$ and $\|x_{t,i}\|_{V_{t-1,i}^{-1}}$ is deterministic given the history to conclude that

$$\begin{aligned} & \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} [Q_{t,i} \beta_i^t(x_{t,i}, \delta)] \\ & \leq \alpha_\delta \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} [Q_{t,i} \min \{1, \|x_{t,i}\|_{V_{t-1,i}^{-1}}\}] \\ & = \alpha_\delta \sum_{i=1}^n \mathbb{E} \left[\sum_{t=1}^T I_{t,i} \min \{1, \|x_{t,i}\|_{V_{t-1,i}^{-1}}\} \right] \\ & = \alpha_\delta \sum_{i=1}^n \mathbb{E} \left[\sum_{j=1}^{m_{T,i}} \min \left\{ 1, \|x_{t_i(j),i}\|_{V_{t_i(j)-1,i}^{-1}} \right\} \right] \\ & \leq \alpha_\delta \sum_{i=1}^n \mathbb{E} \left[\sqrt{m_{T,i} \sum_{j=1}^{m_{T,i}} \min \left\{ 1, \|x_{t_i(j),i}\|_{V_{t_i(j)-1,i}^{-1}}^2 \right\}} \right] \\ & \leq \mathcal{O} \left(\alpha_\delta \cdot n \sqrt{dT \log(1 + T/d)} \right). \end{aligned}$$

Combining the bounds above completes the proof of [Theorem 4.2](#). □

D.3 Proof of [Lemma D.2.2](#)

We follow a similar analysis of the non-contextual case to focus on the regret bound at each round and then sum them up together. Let us consider a round t . For Pandora’s Box problem, we assume

W.L.O.G. that $\sigma_{t,1} \geq \sigma_{t,2} \geq \dots \geq \sigma_{t,n}$. Similar to the non-contextual analysis, we define for all t, i

$$\mathcal{H}_{t,0} = \hat{\mathcal{E}}_t, \quad \text{and} \quad \mathcal{H}_{t,i} = \left(D_{t,1}, \dots, D_{t,i}, \hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n} \right).$$

We have $\mathcal{H}_{t,n} = D_t$ based on this definition. As [Lemma 4.1](#) ensures the stochastic dominance and both problems enjoy a certain monotonicity, we follow the same analysis in [Section 3.2](#) to show that

$$\begin{aligned} \text{Reg}(t) &= \mathbb{E} [R(\sigma_D; D) - R(\sigma_t; D)] \\ &\leq \sum_{i=1}^n \mathbb{E} \left[\underbrace{Q_{t,i} \mathbb{E}_{x \sim \mathcal{H}_{t,i-1}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{y \sim \mathcal{H}_{t,i}} [R(\sigma_t; y) \mid E_{\sigma_t,i}]}_{\text{Term}_{t,i}} \right]. \end{aligned}$$

Now, we define

- $\hat{\mathcal{H}}_{t,i} := \left(D_{t,1}, \dots, D_{t,i-1}, \hat{D}_{t,i}, \hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n} \right)$, and
- $\tilde{\mathcal{H}}_{t,i} := \left(D_{t,1}, \dots, D_{t,i-1}, \mathcal{E}_{t,i}, \hat{\mathcal{E}}_{t,i+1}, \dots, \hat{\mathcal{E}}_{t,n} \right)$.

Based on the definitions of $\mathcal{H}_{t,i}$, $\hat{\mathcal{H}}_{t,i}$, $\tilde{\mathcal{H}}_{t,i}$, we compute the following decomposition:

$$\begin{aligned} \text{Term}_{t,i} &= \mathbb{E}_{x \sim \mathcal{H}_{t,i-1}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{y \sim \mathcal{H}_{t,i}} [R(\sigma_t; y) \mid E_{\sigma_t,i}] \\ &= \underbrace{\mathbb{E}_{x \sim \mathcal{H}_{t,i-1}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{x \sim \tilde{\mathcal{H}}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}]}_{\text{(I)}} \\ &\quad + \underbrace{\mathbb{E}_{x \sim \tilde{\mathcal{H}}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{x \sim \hat{\mathcal{H}}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}]}_{\text{(II)}} \\ &\quad + \underbrace{\mathbb{E}_{x \sim \hat{\mathcal{H}}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{x \sim \mathcal{H}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}]}_{\text{(III)}}. \end{aligned}$$

In the remaining analysis, we retain the same definition of $\tilde{R}_i(\sigma_t; z)$ as in [Eq. \(6\)](#). The analysis conditions on this history before round t , and thus D_t is a product distribution.

Bounding (I). By using the fact that $E_{\sigma_t,i}$ is independent of i -th box's reward, we can bound

$$\begin{aligned} \text{(I)} &= \mathbb{E}_{x \sim \mathcal{H}_{t,i-1}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{x \sim \tilde{\mathcal{H}}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] \\ &= \mathbb{E}_{z \sim \hat{\mathcal{E}}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim \mathcal{E}_{t,i}} [\tilde{R}_i(\sigma_t; z)] \\ &= \int_0^1 (1 - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz - \int_0^1 (1 - F_{\mathcal{E}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz \\ &= \int_0^1 (F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz \\ &\leq \int_0^1 |F_{\mathcal{E}_{t,i}}(z) - F_{\hat{\mathcal{E}}_{t,i}}(z)| dz \\ &\leq \mathcal{O} \left(\sqrt{\frac{L}{m_{t,i}}} \right), \end{aligned}$$

where the first inequality follows from the fact that $\tilde{R}_i(\sigma_t; z)$ is 1-Lipschitz for both Pandora's Box and Prophet Inequality problems (see [Lemma 3.2.1](#) and [Lemma C.3.2](#), respectively), and the last inequality uses the construction of $\hat{\mathcal{E}}_{t,i}$ in [Eq. \(25\)](#).

Bounding (II). One can show that

$$\begin{aligned}
(\text{II}) &= \mathbb{E}_{x \sim \tilde{\mathcal{H}}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{x \sim \mathcal{H}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] \\
&= \mathbb{E}_{z \sim \mathcal{E}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim \hat{D}_{t,i}} [\tilde{R}_i(\sigma_t; z)] \\
&= \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \tilde{R}_i(\sigma_t; \hat{z}_{t_i(j),i}^t) - \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} \tilde{R}_i(\sigma_t; z_{t_i(j),i}^t) \\
&\leq \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} |z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t|,
\end{aligned}$$

where the inequality again the fact that $\tilde{R}_i(\sigma_t; z)$ is 1-Lipschitz for both Pandora's Box and Prophet Inequality problems.

Bounding (III). Then, we have

$$\begin{aligned}
(\text{III}) &= \mathbb{E}_{x \sim \tilde{\mathcal{H}}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] - \mathbb{E}_{x \sim \mathcal{H}_{t,i}} [R(\sigma_t; x) \mid E_{\sigma_t,i}] \\
&= \mathbb{E}_{z \sim \hat{D}_{t,i}} [\tilde{R}_i(\sigma_t; z)] - \mathbb{E}_{z \sim D_{t,i}} [\tilde{R}_i(\sigma_t; z)] \\
&= \int_0^1 (F_{D_{t,i}}(z) - F_{\hat{D}_{t,i}}(z)) \frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) dz \\
&\leq \int_0^1 |F_{D_{t,i}}(z) - F_{\hat{D}_{t,i}}(z)| dz \\
&= \int_0^1 |F_{D_i}(z - \mu_{t,i}) - F_{\tilde{D}_{t,i}}(z - \mu_{t,i})| dz \\
&\leq \mathcal{O}\left(\sqrt{\frac{L}{m_{t,i}}}\right),
\end{aligned}$$

where the first inequality follows from the fact that $\frac{\partial}{\partial z} \tilde{R}_i(\sigma_t; z) \in [-1, 1]$ due to 1-Lipschitzness, and the last inequality follows from the concentration bound given by $\mathcal{E}$.

Putting together. Combining all bounds above, we have

$$\text{Reg}(t) \leq \sum_{i=1}^n \mathcal{O}\left(\mathbb{E}\left[Q_{t,i} \left(\sqrt{\frac{L}{m_{t,i}}} + \frac{1}{m_{t,i}} \sum_{j=1}^{m_{t,i}} |z_{t_i(j),i}^t - \hat{z}_{t_i(j),i}^t|\right)\right]\right).$$

Summing over all $t \in [T]$, we complete the proof of [Lemma D.2.2](#).