
Optimal Contextual Bandits with Knapsacks under Realizability via Regression Oracles

Yuxuan Han^{*†}

Jialin Zeng^{*†}

Yang Wang^{†‡}

Yang Xiang^{†§}

Jiheng Zhang[‡]

[†] Department of Mathematics, HKUST

[‡] Department of Industrial Engineering and Decision Analytics, HKUST

[§] HKUST Shenzhen-Hong Kong Collaborative Innovation Research Institute

Abstract

We study the stochastic contextual bandit with knapsacks (CBwK) problem, where each action, taken upon a context, not only leads to a random reward but also costs a random resource consumption in a vector form. The challenge is to maximize the total reward without violating the budget for each resource. We study this problem under a general realizability setting where the expected reward and expected cost are functions of contexts and actions in some given general function classes \mathcal{F} and \mathcal{G} , respectively. Existing works on CBwK are restricted to the linear function class since they use UCB-type algorithms, which heavily rely on the linear form and thus are difficult to extend to general function classes. Motivated by online regression oracles that have been successfully applied to contextual bandits, we propose the first universal and optimal algorithmic framework for CBwK by reducing it to online regression. We also establish the lower regret bound to show the optimality of our algorithm for a variety of function classes.

1 INTRODUCTION

Contextual Bandits (CB) is a fundamental online learning framework with an exploration-exploitation tradeoff. At each step, the decision maker, to maximize the total reward, takes one out of K actions upon observing the context and then receives a random reward. It has received extensive attention due to its wide range of applications, such as recommendation systems, clinic trials, and

^{*}Equal contribution. Correspondence to Yang Xiang, Jiheng Zhang (<maxiang@ust.hk>, <jiheng@ust.hk>).

Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS) 2023, Valencia, Spain. PMLR: Volume 206. Copyright 2023 by the author(s).

online advertisement (Bietti et al., 2021; Lattimore and Szepesvári, 2020; Slivkins et al., 2019).

However, the standard contextual bandit setting ignores the budget constraint that commonly arises in real-world applications. For example, when a search engine picks an ad to display, it needs to consider both the advertiser’s feature and his remaining budget. Thus, canonical contextual bandits have been generalized to Contextual Bandit with Knapsacks (CBwK) setting, where a selected action at each time t will lead to a context-dependent consumption of several constrained resources and there exists a global budget B on the resources.

In contextual decision-making, it is crucial to model the relationship between the outcome (reward and costs) and the context to design an effective and efficient policy. In a non-budget setting, i.e., without the knapsack constraints, the contextual bandit algorithms fall into two main categories: agnostic and realizability-based approaches. The agnostic approaches aim to find the best policy out of a given policy set Π and make no assumptions on the context-outcome relationships. In this case, one often requires access to a cost-sensitive classification oracle over Π to achieve computational efficiency (Agarwal et al., 2014; Dudik et al., 2011; Langford and Zhang, 2007). On the other hand, realizability-based approaches assume that there exists a given function class \mathcal{F} that models the outcome distribution with the given context. When \mathcal{F} is a linear class, the optimal regret has been achieved by UCB-type algorithms (Abbasi-Yadkori et al., 2011; Chu et al., 2011). Moreover, a unified approach to developing optimal algorithms has recently been studied in Foster et al. (2018), Foster and Rakhlin (2020), Simchi-Levi and Xu (2021) by introducing regression oracles over the class \mathcal{F} . While the realizability assumption could face the problem of model mismatch, it has been shown that empirically realizability-based approaches outperform agnostic-based ones when no model misspecification exists (Foster et al., 2018; Krishnamurthy et al., 2016). However, these oracle-based algorithms only apply to unconstrained contextual bandits without considering the

knapsack constraints.

In the CBwK setting, two similar approaches as above have been explored. For the agnostic approach, Badanidiyuru et al. (2014) and Agrawal et al. (2016) extend the general contextual bandits to the knapsack setting by generalizing the techniques in contextual bandits without knapsack constraints (Agarwal et al., 2014; Dudik et al., 2011). However, as shown in Agrawal and Devanur (2016), the optimization problem using classification oracle in the agnostic setting could be NP-hard even for the linear case. On the other hand, CBwK under linear realizability assumption (i.e., both the reward function class \mathcal{F} and the cost function class \mathcal{G} are linear function classes) and its variants have been studied. Based on the confidence ellipsoid results guaranteed by the linear model, Agrawal and Devanur (2016) develop algorithms for the stochastic linear CBwK with near-optimal regret bounds. Sivakumar et al. (2022) study the linear CBwK under the smoothed contextual setting where the contexts are perturbed by Gaussian noise. However, to our knowledge, no work has considered realizability beyond the linear case. The difficulty lies in the lack of confidence ellipsoid results for the general function class as in linear model assumptions. In this paper, we try to address the following open problem:

Is there a unified algorithm that applies to CBwK under the general realizability assumption?

We answer this question positively by proposing **SquareCBwK**, the first general and optimal algorithmic framework for stochastic CBwK based on online regression oracles. Compared to the contextual bandits, additional resource constraints in CBwK lead to more intricate couplings across contexts, which makes it substantially more challenging to strike a balance between exploration and exploitation. It is thus nontrivial to apply regression oracle techniques to CBwK. The key challenge lies in designing a proper score of actions for decision-making that can well balance the reward and the cost gained over time to maximize the rewards while ensuring no resource is run out.

1.1 Our contributions

In this paper, we address the above issue and successfully apply regression oracles to CBwK under the realizability assumption. Our contributions are summarized in the following three aspects.

Algorithm for CBwK with Online Regression Oracles: In Section 3.1, we propose the first unified algorithm SquareCBwK for solving CBwK under the general realizability assumption, which addresses the above open problem. Motivated by the works in contextual bandits that apply regression oracles to estimate the reward (to which we refer as the reward oracle), we propose to

apply a new cost oracle in SquareCBwK to tackle the additional resource constraints in CBwK. Specifically, under the setting that the budget B scales linearly with the time horizon T , we construct a penalized version of the unconstrained score based on the cost oracle and adjust the penalization adaptive to the resource consumption over time. Our algorithm is able to be instantiated with any available efficient oracle for any general function class to obtain an upper regret bound for CBwK. In this sense, we provide a meta-algorithm that reduces CBwK to regression problems, which is a well-studied area in machine learning with various computationally efficient algorithms for different function classes. In Section 3.3, we show that through different instantiations, SquareCBwK can achieve the optimal regret bound in the linear CBwK and derive new guarantees for more general function classes like non-parametric classes.

Lower Bound Results: In Section 3.2, we develop a new approach for establishing lower bounds for CBwK under the general realizability assumption. Previous works in CBwK under the realizability assumption usually show the optimality of their results by matching the lower bound in the unconstrained setting. However, when the cost function class \mathcal{G} is more complicated than reward function class \mathcal{F} , the lower bound in the unconstrained setting will be very loose since it does not consider the information of \mathcal{G} . Our method, in contrast, can provide lower bounds that include information from both \mathcal{F} and \mathcal{G} . Moreover, we apply this new method to construct lower bounds that match the regret bound given by SquareCBwK for various choices of \mathcal{F} and \mathcal{G} , which demonstrates SquareCBwK’s optimality for different function classes.

Relaxed Assumption on the Budget: While we present our main result under the regime that the budget B scales linearly with the time horizon T , such an assumption may be restrictive compared with previous works. For example, in the linear CBwK, both Sivakumar et al. (2022) and Agrawal and Devanur (2016) allow a relaxed budget $B = \Omega(T^{3/4})$. To close this gap, in Section 4, we design a two-stage algorithm based on SquareCBwK that allows a weaker condition on the budget for solving CBwK under the general realizability assumption. In particular, in the linear CBwK setting, our relaxed condition recovers the $B = \Omega(T^{3/4})$ condition.

1.2 Other related works

Bandit with Knapsack and Demand Learning: One area closely related to our work is bandits with knapsacks (BwK), which does not consider contextual information. The non-contextual BwK problem is first investigated in a general formulation by Badanidiyuru et al. (2018) and further generalized to concave reward/convex constraint setting by Agrawal and Devanur (2014a). Both papers

use the idea of Upper-Confidence Bound as in the non-constrained setting (Auer et al., 2002). Ferreira et al. (2018) apply the BwK framework to the network revenue management problem and develop a Thompson-sampling counterpart for BwK. Their assumption on the demand function is further generalized in the recent works (Chen et al., 2022; Miao and Wang, 2021).

Online Optimization with Knapsacks: One closely related area is online optimization with knapsacks, which can be seen as a full-information variant of the BwK problem: after a decision is made, the feedback for all actions is available to the learner. Such problem often leads to solving online linear/convex programming, which has been studied in Balseiro et al. (2022), Jenatton et al. (2016), Agrawal and Devanur (2014b), Mahdavi et al. (2012), Liu and Grigas (2022) and Castiglioni et al. (2022). In the work of Liu and Grigas (2022), they study online contextual decision-making with knapsack constraints. Specifically, they study the continuous action setting by developing an online primal-dual algorithm based on the ‘‘Smart Predict-then-Optimize’’ framework (Elmachtoub and Grigas, 2022) in the unconstrained setting and leave the problem with bandit feedback open. Our work provides a partial answer to this open problem in the finite-armed setting.

Concurrent works in CBwK: During the review period of our paper, two very recent works appeared on the arxiv.org (Ai et al., 2022; Slivkins and Foster, 2022) which also consider the CBwK problem with general function classes.

In independent and concurrent work, Slivkins and Foster (2022) consider the CBwK problem via regression oracles similar to our work. They propose an algorithm based on the perspective of Lagrangian game that results in a slightly different choice of arm selection strategy than ours. They also attain the same regret bound up to scalar as our Theorem 3.1. Nevertheless, the focus of their work is limited to the setting described in section 3 (i.e., the $B = \Omega(T)$ regime), where they provide only an upper regret bound. In contrast, our work demonstrates the optimality by proving lower bound results for $B = \Omega(T)$ regime. Moreover, we derive a two-stage algorithm and corresponding regret guarantees for the $B = o(T)$ regime that requires less stringent budget assumptions.

Another recent work (Ai et al., 2022) study the CBwK problem with general function classes by combining the re-solving heuristic and the distribution estimation techniques. Their result is interesting in that it may work in some function classes without an efficient online oracle. They also achieve logarithmic regret under suitable conditions thus are more problem-dependent, in contrast to the oracle-based approach, as discussed in section 4 of Foster and Rakhlin (2020). However, their framework involves a distribution estimation procedure

that requires additional assumptions about the regularity of the underlying distribution. Consequently, their method’s regret is influenced by the regularity parameters, potentially leading to suboptimal results.

2 PRELIMINARIES

Notations Throughout this paper, $a \lesssim b$ and $a = O(b)$ ($a \gtrsim b$ and $a = \Omega(b)$) means $a \leq Cb$ ($a \geq Cb$) for some absolute constant C . $a = \tilde{O}(b)$ ($a = \tilde{\Omega}(b)$) means $a = O(b \max\{1, \text{polylog}(b)\})$ ($a = \Omega(b \max\{1, \text{polylog}(b)\})$) and $a \asymp b$ means $a = O(b)$ and $b = O(a)$.

2.1 Basic setup

We consider the stochastic CBwK setting. Given the budget $B \in \mathbb{R}^d$ for d different resources and the time horizon T , at each step, the decision maker needs to select an arm $a_t \in [K]$ upon observing a context $x_t \in \mathcal{X}$ drawn i.i.d. from some unknown distribution $P_{\mathcal{X}}$. Then a reward $r_{t,a_t} \in [0, 1]$ and a consumption vector $\mathbf{c}_{t,a_t} \in [0, 1]^d$ is observed. We assume that $r_{t,a}$ and $\mathbf{c}_{t,a}$ are generated i.i.d. from a fixed distribution parameterized by the given x_t and a . The goal is to learn a policy, a mapping from context to action, that maximizes the total reward while ensuring that the consumption of each resource does not exceed the budget.

Without loss of generality we assume $B_1 = B_2 = \dots = B_d = B$. We focus on the regime $B = \Omega(T)$. That is the budget scales linearly with time. (We relax this assumption on the budget in Section 4.) Moreover, we make a standard assumption that the K -th arm is a null arm that generates no reward or consumption of any resource when being pulled.

Similar to the unconstrained setting (Foster et al., 2018; Foster and Rakhlin, 2020; Simchi-Levi and Xu, 2021), we assume access to two classes of functions $\mathcal{F} \subset (\mathcal{X} \times [K]) \rightarrow [0, 1]$ and $\mathcal{G} \subset (\mathcal{X} \times [K]) \rightarrow [0, 1]^d$ that characterize the expectation of reward and consumption distributions respectively. Note that only one function class \mathcal{F} for the reward distribution is considered in the unconstrained setting, while to fit CBwK, we add another regression class \mathcal{G} to model the consumption distribution. We assume the following realizability condition.

Assumption 2.1. *There exists some $f^* \in \mathcal{F}$ and $\mathbf{g}^* \in \mathcal{G}$ such that $f^*(x, a) = \mathbb{E}[r_{t,a}|x_t = x]$ and $\mathbf{g}^*(x, a) = \mathbb{E}[\mathbf{c}_{t,a}|x_t = x], \forall a \in [K]$.*

The algorithm benchmark is the best dynamic policy, which knows the contextual distribution $P_{\mathcal{X}}$, f^* , and \mathbf{g}^* and can dynamically maximize total reward given the historical information and the current context. We denote OPT_{DP} as the expected total reward of the best dynamic policy. The goal of the decision-maker is to minimize the following regret:

Definition 2.1. $\text{Reg}(T) := \text{OPT}_{\text{DP}} - \mathbb{E}[\sum_{t=1}^{\tau} r_{t,a_t}]$, where τ is the stopping time when there exists some $j \in [d]$ s.t. $\sum_{t=1}^{\tau} (\mathbf{c}_{t,a_t})_j > B - 1$.

Due to the intractability of the best dynamic policy in Definition 2.1, we consider a static relaxation problem that provides an upper bound of OPT_{DP} .

Denote Δ^K the set of probability distributions over the action set $[K]$. The following program aims to find the best static randomized policy $p^* : \mathcal{X} \rightarrow \Delta^K$ that maximizes the expected per-round reward while ensuring resources are not exceeded in expectation.

$$\begin{aligned} \max_{p: \mathcal{X} \rightarrow \Delta^K} \mathbb{E}_{x \sim P_{\mathcal{X}}} \left[\sum_{a \in [K]} p_a(x) f^*(x, a) \right] \\ \text{s.t. } \mathbb{E}_{x \sim P_{\mathcal{X}}} \left[\sum_{a \in [K]} p_a(x) \mathbf{g}^*(x, a) \right] \leq B/T \cdot \mathbf{1}. \end{aligned} \quad (1)$$

Lemma 2.1. *Let OPT denote the value of the optimal static policy (1), then we have $T\text{OPT} \geq \text{OPT}_{\text{DP}}$.*

Lemma 2.1 can be derived directly following the proof of Lemma 1 in Agrawal and Devanur (2016), where they consider the linear CBwK case but the reasoning therein is completely independent of the linear structure. With Lemma 2.1, we can control the regret bound by considering $T\text{OPT} - \mathbb{E}[\sum_{t=1}^{\tau} r_{t,a_t}]$.

2.2 Online Regression Oracles

Under the realizability Assumption 2.1, We introduce the online regression problem and the notion of online regression oracles.

The general setting of online regression problem with input space \mathcal{Z} , output space \mathcal{Y} , function class \mathcal{H} and loss $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is described as following: At the beginning of the game, the environment chooses an $h^* : \mathcal{Z} \rightarrow \mathcal{Y}$, $h^* \in \mathcal{H}$ as the underlying response generating function. Then at each round t , (i) the learner receives an input $z_t \in \mathcal{Z}$ possibly chosen in adversarial by the environment, (ii) the learner then predicts a value $\hat{y}(z_t) \in \mathcal{Y}$ based on historical information, (iii) the learner observes a noisy response of $h^*(z_t)$ and suffers a loss $\ell(\hat{y}(z_t), h^*(z_t))$. The goal of the learner is to minimize the cumulative loss

$$\text{Reg}_{\text{sq}}(T; h^*) := \sum_{t=1}^T \ell(\hat{y}(z_t), h^*(z_t)).$$

In our problem, we assume our access to oracles $\mathcal{R}^r, \mathcal{R}^c$ for two online regression problems, respectively:

Reward Regression Oracle The reward regression oracle \mathcal{R}^r is assumed to be an algorithm for the online regression problem with $\mathcal{Z} = \mathcal{X} \times [K], \mathcal{Y} = [0, 1], \mathcal{H} = \mathcal{F}, \ell(y_1, y_2) = (y_1 - y_2)^2$ so that when \hat{y}_t is generated by

\mathcal{R}^r , there exists some Reg_{sq}^r as a function of T such that

$$\text{Reg}_{\text{sq}}(T; f^*) \leq \text{Reg}_{\text{sq}}^r(T), \quad \forall f^* \in \mathcal{F}. \quad (2)$$

Cost Regression Oracle The cost regression oracle \mathcal{R}^c is assumed to be an algorithm for the online regression problem with $\mathcal{Z} = \mathcal{X} \times [K], \mathcal{Y} = [0, 1]^d, \mathcal{H} = \mathcal{G}, \ell(\mathbf{y}_1, \mathbf{y}_2) = \|\mathbf{y}_1 - \mathbf{y}_2\|_{\infty}^2$ so that when $\hat{\mathbf{y}}_t$ is generated by \mathcal{R}^c , there exists some Reg_{sq}^c as a function of T such that

$$\text{Reg}_{\text{sq}}(T; \mathbf{g}^*) \leq \text{Reg}_{\text{sq}}^c(T), \quad \forall \mathbf{g}^* \in \mathcal{G}. \quad (3)$$

2.3 Online Mirror Descent

Our algorithm adopts an online primal-dual framework to tackle the challenge brought by knapsack constraints. Our strategy of updating the dual variable λ_t falls into the general online convex optimization (OCO) framework. In the OCO problem with parameter set Λ , adversary class \mathcal{L} and time horizon T , the learner needs to choose $\lambda_t \in \Lambda$ adaptively at each round t . After each choice, he will observe an adversarial convex loss $L_t \in \mathcal{L}$ and pay the cost $L_t(\lambda_t)$. The goal of the learner is to minimize the cumulative regret:

$$\text{Reg}_{\text{OCO}}(T; \Lambda, \mathcal{L}) := \sum_{t=1}^T L_t(\lambda_t) - \min_{\lambda \in \Lambda} \sum_{t=1}^T L_t(\lambda).$$

In our designed adversary, $L_t(\lambda) = \langle \frac{B}{T} \cdot \mathbf{1} - \mathbf{c}_{t,a_t}, \lambda \rangle$ and minimizing L_t corresponds to penalize the violation of the budget. We focus on the following adversary class and parameter set:

$$\begin{aligned} \mathcal{L} &= \{L(\lambda) := \langle \theta, \lambda \rangle : \theta \in \mathbb{R}^d, \|\theta\|_{\infty} \leq 1\} \\ \Lambda &= \{\lambda \in \mathbb{R}^d : \lambda \geq 0, \|\lambda\|_1 \leq Z\}, \end{aligned}$$

where $Z > 0$ is the ℓ_1 radius of Λ to be determined. The online mirror descent (OMD) algorithm (Hazan et al., 2016; Shalev-Shwartz et al., 2012) is a simple and fast algorithm achieving the optimal OCO regret in the non-euclidean geometry. OMD follows the update rule

$$\lambda_t = \arg \min_{\lambda \in \Lambda} (\nabla L_{t-1}(\lambda_{t-1}), \lambda) + \frac{1}{\eta_t} D_h(\lambda, \lambda_{t-1}), \quad (4)$$

where h is the metric generating function, and D_h is the associated Bregman divergence. After adding a slack variable and re-scaling, the OCO problem over \mathcal{L}, Λ is equivalent to the OCO problem over

$$\begin{aligned} \tilde{\mathcal{L}} &= \{\tilde{L}(\tilde{\lambda}) := \langle \tilde{\theta}, \tilde{\lambda} \rangle : \tilde{\theta} \in \mathbb{R}^{d+1}, \|\tilde{\theta}\|_{\infty} \leq Z, \tilde{\theta}_{d+1} = 0\} \\ \tilde{\Lambda} &= \{\tilde{\lambda} \in \mathbb{R}^{d+1} : \tilde{\lambda} \geq 0, \|\tilde{\lambda}\|_1 = 1\}, \end{aligned}$$

which can be solved via the normalized Exponentiated Gradient (i.e., selecting h as the negative entropy function). In this case, the OMD algorithm has the following OCO regret guarantee (Hazan et al., 2016; Shalev-Shwartz et al., 2012):

Lemma 2.2. Setting $\eta_t = \eta = O(\frac{\log d}{\sqrt{T}})$, the OMD yields regret

$$\text{Reg}_{\text{OCO}}(T; \Lambda, \mathcal{L}) \lesssim Z\sqrt{T \log d} \quad (5)$$

As in the literature (Agrawal and Devanur, 2016; Castiglioni et al., 2022) that introduces the online primal-dual framework, Lemma 2.2 plays an essential role in controlling the regret induced by knapsack constraints in SquareCBwK.

3 ALGORITHM AND THEORETICAL GUARANTEES

We are ready to present our main algorithm SquareCBwK for solving stochastic CBwK under the general realizability assumption. We first present the algorithm and theoretical results under $B = \Omega(T)$ regime for simplicity of algorithm design. Algorithm and analysis for more general choices of B are presented in section 4.

3.1 The SquareCBwK Algorithm

The main body of SquareCBwK has a similar structure as the SquareCB algorithm for contextual bandits (Foster and Rakhlin, 2020), but with significant changes necessary to handle the knapsack constraints. SquareCBwK is presented in Algorithm 1 with three key modules: prediction through oracles, arm selection scheme, and dual update through OMD.

Prediction through oracles At each step, after observing the context, SquareCBwK will simultaneously access the reward oracle \mathcal{R}^r and cost oracle \mathcal{R}^c to predict the reward and cost of each action for this round. Then these two predicted scores are incorporated through the following Lagrangian to form a final predicted score $\hat{\ell}_t$:

$$\hat{\ell}_{t,a} := \hat{r}_{t,a} + \lambda_t^T (\mathbf{1} \cdot B/T - \hat{c}_{t,a}), \forall a \in [K].$$

Arm selection scheme After computing the predicted score $\hat{\ell}_t$, we employ the probability selection strategy in Abe and Long (1999): We choose the greedy action score evaluated by $\hat{\ell}_t$ as the benchmark and select each arm a with the probability $p_{t,a}$ that is inversely proportional to the gap between the arm's score and the benchmark. This strategy strikes a balance between exploration and exploitation: When the predicted score for an action is close to the greedy action, we tend to explore it with the probability roughly as $1/K$, otherwise with a very small chance.

Dual update through OMD After choosing the arm a_t , the new data $\{x_t, a_t, r_{t,a_t}, c_{t,a_t}\}$ will be fed into \mathcal{R}^c , \mathcal{R}^d and OMD. We then update the dual variable $\lambda_t \in \Lambda$ successively through OMD.

Algorithm 1: SquareCBwK

Input: Time horizon T , total budget initial B , learning rate $\gamma > 0$, online regression oracle for reward \mathcal{R}^r and cost \mathcal{R}^c , radius $Z = T/B$ of the parameter set Λ .

```

1 Initialization: initialize  $\mathcal{R}^r$ ,  $\mathcal{R}^c$  and  $\mathcal{R}^d$ .
2 for  $t = 1, \dots, T$  do
3   Observe context  $x_t$ .
4    $\mathcal{R}^r$  predicts  $\hat{r}_{t,a}$ ,  $\mathcal{R}^c$  predicts  $\hat{c}_{t,a}$ ,  $\forall a \in [K]$ 
5   Compute  $\hat{\ell}_{t,a} := \hat{r}_{t,a} + \lambda_t^T (B/T \cdot \mathbf{1} - \hat{c}_{t,a})$ ,
      $\forall a \in [K]$ .
6   Let  $b_t = \arg \max_{a \in [K]} \hat{\ell}_{t,a}$ .
7   For each  $a \neq b_t$ , define
     
$$p_{t,a} := \frac{1}{K + \gamma(\hat{\ell}_{t,b_t} - \hat{\ell}_{t,a})}$$
 and let
     
$$p_{t,b_t} = 1 - \sum_{a \neq b_t} p_{t,a}$$
.
8   Sample arm  $a_t \sim p_t$ . Observe reward  $r_{t,a_t}$  and
     cost  $c_{t,a_t}$ .
9   if  $\exists j \in [d], \sum_{t'=1}^t c_{t',a_t}[j] \geq B - 1$  then
10    | Exit
11  end
12  Feed  $\mathcal{R}^r$  with data  $\{(x_t, a_t), r_{t,a_t}\}$  and  $\mathcal{R}^c$  with
     data  $\{(x_t, a_t), c_{t,a_t}\}$ .
13  Feed OMD with  $c_{t,a_t}$  and OMD updates
      $\lambda_{t+1} \in \Lambda$ .
14 end

```

Compared with SquareCB in Foster and Rakhlin (2020), we introduce three novel technical elements to deal with knapsack constraints: First, we apply a new cost oracle to generate the cost prediction. Next, to balance the reward and the cost prediction over time, we propose the predicted Lagrangian to construct a proper score function $\hat{\ell}_t$. Finally, we introduce OMD to update the dual variable so that the predicted scores $\hat{\ell}_t$ adapt to the resource consumption throughout the process. Notably, the l_1 radius of the parameter set Λ is carefully chosen to be T/B since we expect λ_t can capture the sensitivity of the optimal static policy (1) to knapsack constraints violations over time. Specifically, if we increase the budget B by ε , the increased reward over T rounds is at most $\frac{T \text{OPT}}{B} \varepsilon$. This observation suggests that Z , the radius of Λ should be at least $\frac{T \text{OPT}}{B}$. On the other hand, Z should be of constant level so that OMD achieves optimal regret bounds $O(\sqrt{T})$ by Lemma 2.2. Therefore, setting $Z = \frac{T \text{OPT}}{B}$ will be the desired choice. In practice, since OPT is unknown, we need to estimate Z so that $\frac{T \text{OPT}}{B} \leq Z \lesssim \frac{T \text{OPT}}{B}$. The regime $B = \Omega(T)$ guarantees that $\frac{T}{B}$ is approximately $\frac{T \text{OPT}}{B}$, without the need to further estimate OPT. This is why we set $Z = \frac{T}{B}$ in SquareCBwK. We will discuss estimating OPT when $B = \Omega(T)$ fails to hold in Section 4.

Now we state the theoretical guarantee of SquareCBwK:

Theorem 3.1. *Considering the regime $B = \Omega(T)$ under Assumption 2.1, if the output of \mathcal{R}^r and \mathcal{R}^c satisfy (2) and (3), respectively, denote*

$$\gamma = \sqrt{KT / (\text{Reg}_{\text{Sq}}^r(T) + (T/B + 1)^2 \text{Reg}_{\text{Sq}}^c(T) + 4\log(2T))},$$

then SquareCBwK achieves the regret

$$\begin{aligned} \text{Reg}(T) \lesssim & \sqrt{KT \cdot (\text{Reg}_{\text{Sq}}^r(T) + \log(dT))} \\ & + \left(\frac{T}{B} + 1\right) \sqrt{KT \text{Reg}_{\text{Sq}}^c(T)} \end{aligned}$$

Compared to the result in Foster and Rakhlin (2020) for the unconstrained contextual bandit problem, our regret bound has an additional dependency on $\text{Reg}_{\text{sq}}^c(T)$, which is a natural outcome under the budget-setting. Moreover, with knapsack constraints, the optimal static policy will be a distribution over actions rather than pulling a single optimal arm. As a result, in the proof of Theorem 3.1, we need to adapt the argument related to the probability selection strategy in Foster and Rakhlin (2020) to this significant change and derive a lower bound for the total expected predicted scores. Then we split the total expected reward from the expected Lagrangian scores and control the regret incurred by the early stopping time $\tau < T$ using the regret of OMD and special radius selection Z . We relate expectation with realization to obtain the final regret bound.

Theorem 3.1 provides an upper bound for the regret of stochastic CBwK by reducing it to regression, a basic supervised learning task. We further show such a reduction is optimal for various function classes in Section 3.2 where we design a novel way to derive lower bounds for the regret of CBwK that match the upper bound given by Theorem 3.1. With this optimal reduction, our framework is quite general and flexible: it can be instantiated with any available efficient and optimal oracles for general \mathcal{F} and \mathcal{G} (we also allow \mathcal{G} to be different from \mathcal{F}), then the optimal regret for CBwK can be directly given by Theorem 3.1.

3.2 Lower Bound Results

To demonstrate the optimality of Theorem 3.1, we need to discuss whether the dependency on Reg_{sq}^r and Reg_{sq}^c is tight. In the unconstrained setting, Foster and Rakhlin (2020) shows the tightness result of Reg_{sq}^r for a wide range of nonparametric classes \mathcal{F} . Since the contextual bandit problem can be seen as a special case of CBwK problem with $B = T$ and $d = 1$, the optimality results in Foster and Rakhlin (2020) can be utilized to show the tight dependency on Reg_{sq}^r in Theorem 3.1, which also indicates the tight dependency on Reg_{sq}^c when \mathcal{G} shares the same structure as \mathcal{F} since the term Reg_{sq}^c can be absorbed by Reg_{sq}^r in this case. However, when the complexity of \mathcal{G} is much higher than \mathcal{F} , the lower regret bound in the unconstrained setting will be loose compared with the

upper bound in Theorem 3.1. To close this gap, we obtain a general result that establishes lower bounds for CBwK concerning cost function class \mathcal{G} based on its in-separation property.

To present our result, for a general function class \mathcal{H} , we first introduce a new concept that characterizes the difficulty of the **unconstrained** contextual bandit problems with \mathcal{H} as the reward function class:

Definition 3.1 (α -inseparable class). *For a fixed time horizon T , we say a given function class \mathcal{H} of functions from $\mathcal{X} \times \{0, 1\}$ to $[0, 1]$ is α -inseparable with respect to some $P_{\mathcal{X}}$ over \mathcal{X} , if there exist $\mathcal{U}_{\alpha} \subset \mathcal{H}$ and an absolute constant $c > 0$ independent of T such that*

1. *For all $h \in \mathcal{U}_{\alpha}$, denoting $a^*(x; h) = \underset{a \in \{0, 1\}}{\text{argmax}} h(x, a)$, it holds that*

$$\begin{aligned} & \mathbb{E}_{P_{\mathcal{X}}}[h(x, a^*(x; h))] - \alpha \\ & \geq \mathbb{E}_{P_{\mathcal{X}}}[h(x, 1 - a^*(x; h))] \geq \frac{1}{4} \end{aligned}$$

2. *For every dynamic policy π , there exist some $h \in \mathcal{U}_{\alpha}$ and a distribution of $r_{t,a}$ with $\mathbb{E}[r_{t,a}|x_t = x] = h(x, a)$, such that*

$$\mathbb{E}_{x,r} \left[\sum_{t=1}^T \mathbf{1}\{\pi_t(x_t) \neq a^*(x_t; h)\} \right] > cT.$$

We refer to the above properties 1 and 2 as the α -inseparable property since they indicate that even the optimal action outperforms the sub-optimal action with a reward gap larger than α , there exists no dynamic policy that can distinguish the optimal action without exploring at least $\Omega(T)$ steps. One straightforward observation from Definition 3.1 is that, if \mathcal{H} is α -inseparable, then

$$\text{Reg}_{\text{CB}}(T; \mathcal{H}) \gtrsim \alpha T,$$

where $\text{Reg}_{\text{CB}}(T; \mathcal{H})$ is the optimal minimax regret bound of unconstrained CB problem with \mathcal{H} as the expected reward class. Here the minimum is taken over all possible dynamic policies and the maximum is taken over all possible expected reward function $h^* \in \mathcal{H}$, all $P_{\mathcal{X}}$ over \mathcal{X} and all conditional distributions P_r of rewards with $E[r_{t,a}|x] = h^*(x, a)$.

Indeed, the α -inseparable property has been applied implicitly for many classes in previous works to derive tight lower bounds for $\text{Reg}_{\text{CB}}(T; \mathcal{H})$, e.g., linear class (Chu et al., 2011), Hölder class (Rigollet and Zeevi, 2010) and general nonparametric class (Foster and Rakhlin, 2020).

Now we state our main theorem that establishes regret lower bounds for CBwK with α -inseparable classes:

Theorem 3.2. *Consider the class of CBwK problems with two non-null arms, $d = 1$, the reward class \mathcal{F} and the cost class \mathcal{G} , where \mathcal{G} is α -inseparable with respect to*

some P_0 over \mathcal{X} , and there exists some $f_0 \in \mathcal{F}$ such that $f_0(x, 0) = f_0(x, 1) = \Omega(1)$ a.s. under P_0 . Then there exists T and $B^* > T/8$ such that for every dynamic policy π , there exists $f^* \in \mathcal{F}, g^* \in \mathcal{G}$ and distribution $P_{\mathcal{X}}^*$ of context, distribution P_r^*, P_c^* on rewards and costs with $\mathbb{E}[r_{t,a}|x_t] = f^*(x_t, a), \mathbb{E}[c_{t,a}|x_t] = g^*(x_t, a)$ such that when one runs π on the CBwK instance with budget B^* , context distribution $P_{\mathcal{X}}^*$, reward distribution P_r^* and cost distribution P_c^* , it holds that

$$\text{Reg}_{\pi}(T) \gtrsim \max\{\text{Reg}_{\text{CB}}(T; \mathcal{F}), \alpha T\}.$$

The condition on the existence of f_0 is a technical assumption that ensures one can always construct an instance with the achieved reward independent of the action taken. For such instance, the regret of a policy is only determined by its over-cost of the resource, which can be lower bounded by αT . As discussed above, the αT term is a lower bound of $\text{Reg}_{\text{CB}}(T; \mathcal{G})$, which is the CB lower bound when \mathcal{G} is the reward class. In this sense, Theorem 3.2 reveals that the lower bound of a CBwK problem with reward class \mathcal{F} and cost class \mathcal{G} can be established by just considering unconstrained CB regret bounds with reward classes \mathcal{F} or \mathcal{G} .

3.3 Applications

In this section, we instantiate SquareCBwK with different oracles for the generalized linear class and nonparametric function classes, respectively. We show that Theorem 3.1 and Theorem 3.2 can provide tight regret upper and lower bounds for these classes. The regret bounds and selection of oracles are summarized in Table 1. As far as we know, no previous results study CBwK beyond the linear setting or consider the tightness of lower bounds when \mathcal{F} and \mathcal{G} are different even for the linear case. Although here we focus on the generalized linear and nonparametric classes, SquareCBwK can also be instantiated with available oracles of other function classes, e.g., the kernel classes and uniform convex Banach spaces discussed in section 2.3 of Foster and Rakhlin (2020).

For the examples considered in this section, the vector-valued function class \mathcal{G} is a product of d same $[0, 1]$ -valued function class $\tilde{\mathcal{G}}$, i.e.,

$$\mathcal{G} = \tilde{\mathcal{G}}^d = \{\mathbf{g} : \mathbf{g} = (g_1, \dots, g_d), g_i \in \tilde{\mathcal{G}}, \forall i \in [d]\}.$$

In this case, we can construct the online regression oracle \mathcal{R}^c satisfying (3) with $\text{Reg}_{\text{sq}}^c(T) \leq d\widetilde{\text{Reg}}_{\text{sq}}^c(T)$ from any oracle $\tilde{\mathcal{R}}^c$ over $\tilde{\mathcal{G}}$ satisfying (3) with $\widetilde{\text{Reg}}_{\text{sq}}^c(T)$. We provide such construction in Appendix A.3. Here we only specify $\tilde{\mathcal{G}}$ and $\tilde{\mathcal{R}}^c$ in the following examples.

Generalized Linear CBwK In the generalized linear CBwK setting, there exist known feature maps $\phi_1 : \mathcal{X} \times [K] \rightarrow \mathbb{B}^{m_1}, \phi_2 : \mathcal{X} \times [K] \rightarrow \mathbb{B}^{m_2}$, where \mathbb{B}^m is the

unit ℓ_2 ball in \mathbb{R}^m , and link functions $\sigma_i : [-1, 1] \rightarrow [0, 1]$ satisfying $0 < c \leq \sigma_i'(x) < 1$ for $i = 1, 2$. The function classes are selected as

$$\begin{aligned} \mathcal{F} &= \{(x, a) \rightarrow \sigma_1(\langle \theta, \phi_1(x, a) \rangle), \theta \in \mathbb{B}^{m_1}\}, \\ \tilde{\mathcal{G}} &= \{(x, a) \rightarrow \sigma_2(\langle \theta, \phi_2(x, a) \rangle), \theta \in \mathbb{B}^{m_2}\}. \end{aligned}$$

In this case, selecting $\tilde{\mathcal{R}}^r, \tilde{\mathcal{R}}^c$ as the Newtonized GLMtron oracle (Foster and Rakhlin (2020), Proposition 3) achieves $\text{Reg}_{\text{sq}}^r(T) \lesssim m_1 \log(T), \text{Reg}_{\text{sq}}^c(T) \lesssim dm_2 \log(T)$. Then Theorem 3.1 implies the $O(\sqrt{(m_1 + dm_2)KT})$ regret of SquareCBwK. On the other hand, we can verify $\tilde{\mathcal{G}}$ is $\sqrt{m_2/T}$ -inseparable (see Appendix A.3), and $\text{Reg}_{\text{CB}}(T; \mathcal{F}) \gtrsim \sqrt{m_1 T}$. Then applying Theorem 3.2 leads to a $\Omega(\sqrt{(m_1 + dm_2)KT})$ lower bound when $d = 1$. Our upper bound and lower bounds imply the optimality of SquareCBwK with respect to T and m_1, m_2 . Moreover, besides the $O(\sqrt{(m_1 + m_2d)KT})$ achieved by the GLMtron oracle, we can also select the Online Gradient Descent oracle (Foster and Rakhlin (2020), Proposition 2) for SquareCBwK to get $\tilde{O}(\sqrt{KdT}^{3/4})$ regret, which has worse dependency on T but is independent of dimension.

By selecting σ as the identical map, and assuming in addition the ranges of ϕ_1, ϕ_2 lie in $\mathbb{R}_+^{m_1}, \mathbb{R}_+^{m_2}$, the generalized linear class covers the linear model as a special case. Thus our algorithm also applies to linear CBwK. Compared with the $\tilde{O}((m_1 + m_2)\sqrt{T})$ regret achieved in Agrawal and Devanur (2016), SquareCBwK with Newtonized GLMtron oracle has an additional dependency on K, d , but improves the dependency of m_1, m_2 to $\sqrt{m_1}, \sqrt{m_2}$, which matches the established lower bound. In Appendix C, we perform simulations for linear CBwK, comparing the dependencies on time horizon T , dimension m , and number of arms K of SquareCBwK utilizing Newtonized GLMtron and Online Gradient Descent oracles with those of the LinUCB in Agrawal and Devanur (2016). These simulations provide numerical verification of the aforementioned theoretical guarantees.

CBwK with nonparametric function classes We say a function class \mathcal{W} is in the nonparametric regime if its metric entropy $H(\mathcal{W}, \varepsilon)$ scales in ε^{-p} for some $p > 0$ (Rakhlin et al., 2017). More precisely, we say a class \mathcal{W} is p -nonparametric if $H(\mathcal{W}, \varepsilon) \lesssim \varepsilon^{-p}, \forall \varepsilon > 0$. We make the additional assumption that both \mathcal{F} and \mathcal{G} tensorizes: There exists p_1 -nonparametric and p_2 -nonparametric classes \mathcal{W}, \mathcal{V} so that

$$\begin{aligned} \mathcal{F} &= \{f : (x, a) \rightarrow w_a(x), w \in \mathcal{W}\}, \\ \mathcal{G} &= \{\mathbf{g} : (x, a) \rightarrow \mathbf{v}_a(x), \mathbf{v} \in \mathcal{V}\}. \end{aligned}$$

The p -nonparametric class is general enough to cover many classes including the Hölder class considered in most nonparametric CB literature (Hu et al., 2020; Rigollet and Zeevi, 2010; Slivkins, 2011). In this case, an Vovk's aggregation based oracle is proposed in Foster

Table 1: Regret bound results of SquareCBwK with different instantiations when $B = \Omega(T)$.

\mathcal{F}	$\tilde{\mathcal{G}}$	\mathcal{R}^r and $\tilde{\mathcal{R}}^c$	Previous Results	Regret by Thm 3.1	Lower Bound by Thm 3.2
m_1 -dim Linear	m_2 -dim Linear	Newtonized GLMtron*	$\tilde{O}((m_1 + m_2)\sqrt{T})^\dagger$	$\tilde{O}(\sqrt{(m_1 + dm_2)KT})$	$\Omega(\sqrt{(m_1 + m_2)T})$
m_1 -dim Generalized Linear	m_2 -dim Generalized Linear		N.A.		
p_1 -Nonparametric	p_2 -Nonparametric	Vovk's Aggregation [§]		$\tilde{O}((KT)^{\frac{1+p_1}{2+p_1}} + \sqrt{d}(KT)^{\frac{1+p_2}{2+p_2}})$	$\tilde{\Omega}(T^{\frac{1+p_1}{2+p_1}} + T^{\frac{1+p_2}{2+p_2}})$

* Proposition 3, Foster and Rakhlin (2020) ; [§] Theorem 3, Foster and Rakhlin (2020) ; [†] Theorem 2, Agrawal and Devanur (2016) (while Agrawal and Devanur (2016) study the setting $m_1 = m_2$, it is straightforward to extend their result to general m_1, m_2 .)

and Rakhlin (2020) for \mathcal{F} and $\tilde{\mathcal{G}}$ with $\text{Reg}_{\text{sq}}^r(T) = \tilde{O}((KT)^{1-\frac{2}{2+p_1}})$, $\text{Reg}_{\text{sq}}^c(\mathcal{G}) = \tilde{O}(d(KT)^{1-\frac{2}{2+p_2}})$, thus Theorem 3.1 implies the $\tilde{O}((KT)^{\frac{1+p_1}{2+p_1}} + \sqrt{d}(KT)^{\frac{1+p_2}{2+p_2}})$ regret guarantee of SquareCBwK. By verifying that $\tilde{\mathcal{G}}$ is $T^{1-\frac{2}{2+p_2}}$ -inseparable (see Appendix A.3) and applying Theorem 3.2, we also establish a tight lower bound $\tilde{\Omega}(T^{\frac{1+p_1}{2+p_1}} + T^{\frac{1+p_2}{2+p_2}})$. Our upper and lower bound results imply the *universality* of the SquareCBwK: for every general nonparametric \mathcal{F} and $\tilde{\mathcal{G}}$ there always exist choices of \mathcal{R}^r and $\tilde{\mathcal{R}}^c$ such that SquareCBwK achieves the optimal regret bound with respect to T and the complexity parameters p_1, p_2 of $\mathcal{F}, \tilde{\mathcal{G}}$.

4 ALGORITHM WITH RELAXED ASSUMPTION ON B

In this section, we aim to relax the assumption $B = \Omega(T)$. When $B = o(T)$, applying the Theorem 3.1 in this scenario will lead to a sub-optimal result due to its dependency on the $\frac{T}{B}$ factor. Indeed, the true dependency should be $\frac{T \text{OPT}}{B}$ as we discussed before. We can get rid of this $\frac{T}{B}$ factor by replacing the radius of Λ with a factor Z that approximates $\frac{T \text{OPT}}{B}$ better. We propose a two-stage algorithm, in which $\frac{T \text{OPT}}{B}$ is approximated by Z in the first phase, and in the second phase Algorithm 1 is run with Z .

While such a two-stage design has been used in the linear setting (Agrawal and Devanur, 2016; Sivakumar et al., 2022), their designs depend on the special structure of linear classes and the self-normalized martingale concentration, which cannot be extended to general classes. To estimate Z for more general \mathcal{F}, \mathcal{G} , we introduce the statistical regression oracles $\mathcal{R}_{\text{est}}^r, \mathcal{R}_{\text{est}}^c$ with the following assumptions:

Assumption 4.1. For any a , given a dataset \mathcal{D}_a containing M i.i.d. samples $\{(x_i, r_{i,a}, c_{i,a})\}_{i=1}^M \sim P_{\mathcal{X}}$, the output \hat{f}, \hat{g} of $\mathcal{R}_{\text{est}}^r, \mathcal{R}_{\text{est}}^c$ with input \mathcal{D}_a satisfy

$$\mathbb{E}_{x \sim P_{\mathcal{X}}} [(\hat{f}(x, a) - f(x, a))^2] \leq \mathcal{E}_\delta(M; \mathcal{F}),$$

$$\mathbb{E}_{x \sim P_{\mathcal{X}}} [\|\hat{g}(x, a) - \mathbf{g}(x, a)\|_\infty^2] \leq \mathcal{E}_\delta(M; \mathcal{G}).$$

with probability at least $1 - \delta$.

Indeed, as long as there exists an online regression oracle satisfying (2) and (3), we can apply the standard

Algorithm 2: TwoStage-SquareCBwK

Input: T, B, T_0 , estimate oracles $\mathcal{R}_{\text{est}}^r, \mathcal{R}_{\text{est}}^c$, other parameters in Algorithm 1.

- 1 **Initialization:** $\mathcal{D}_{0,a}^r, \mathcal{D}_{0,a}^c = \emptyset, \forall a \in [K]$.
 - 2 **for** $a = 1, \dots, K$ **do**
 - 3 **for** $t = (a-1)T_0 + 1, \dots, aT_0$ **do**
 - 4 Play arm a and observe r_t, c_t .
 - 5 $\mathcal{D}_{0,a}^r = \mathcal{D}_{0,a}^r \cup \{r_t\}$,
 - 6 $\mathcal{D}_{0,a}^c = \mathcal{D}_{0,a}^c \cup \{c_t\}$.
 - 7 **end**
 - 8 Set $\hat{f}_0(\cdot, a)$ the output of $\mathcal{R}_{\text{est}}^r$ with input $\mathcal{D}_{0,a}^r$.
 - 9 Set $\hat{g}_0(\cdot, a)$ the output of $\mathcal{R}_{\text{est}}^c$ with input $\mathcal{D}_{0,a}^c$.
 - 10 **end**
 - 11 **for** $t = KT_0 + 1, \dots, (K+1)T_0$ **do**
 - 12 Pull arm a_t arbitrarily.
 - 13 **end**
 - 14 Solve the linear programming (6) and get $\widehat{\text{OPT}}(T_0)$.
 - 15 Set the radius of Λ to Z as computed in Lemma 4.1.
 - 16 Set remaining resource as $\mathbf{B}' = \mathbf{B} - (K+1)T_0 \mathbf{1}$.
 - 17 Run Algorithm 1 for $T - (K+1)T_0$ rounds with remaining budget \mathbf{B}' and Z .
-

online-to-batch (OTB) method to construct a statistical regression oracle that satisfies $\mathcal{E}_\delta(M; \mathcal{F}) \lesssim \frac{\text{Reg}_{\text{sq}}^r(M) \log(1/\delta)}{M}$ and $\mathcal{E}_\delta(M; \mathcal{G}) \lesssim \frac{\text{Reg}_{\text{sq}}^c(M) \log(1/\delta)}{M}$. We leave the construction of OTB oracle and the proof of such estimation error to Appendix B.1.

The Stage 1 of Algorithm 2 includes the first $(K+1)T_0$ rounds to estimate $\frac{T \text{OPT}}{B}$. For the first KT_0 rounds, we pull each arm evenly for T_0 times regardless of the context and gather outcomes. We then use the oracles over the collected data to generate predictors \hat{f}_0 and \hat{g}_0 . For the next T_0 rounds, we collect the contexts and pull arms arbitrarily. The contextual information in the latter T_0 rounds is used to estimate OPT by solving the following linear programming over $(\Delta^K)^{T_0}$:

$$\begin{aligned} & \max_{p \in (\Delta^K)^{T_0}} \frac{1}{T_0} \sum_{t \in \mathcal{T}_0} \sum_{a \in [K]} p_{t,a} \hat{f}_0(x_t, a) \\ & \text{subject to } \frac{1}{T_0} \sum_{t \in \mathcal{T}_0} \sum_{a \in [K]} p_{t,a} \hat{g}_0(x_t, a) \leq \frac{\mathbf{B}}{T} + 2\mathcal{M}(T_0), \end{aligned} \quad (6)$$

where we denote

$$\mathcal{T}_0 := \{t : KT_0 + 1 \leq t \leq (K+1)T_0\},$$

Table 2: Regret bound results of TwoStage-SquareCBwK with different instantiation when $B = o(T)$.

\mathcal{F} and $\tilde{\mathcal{G}}$	Theorem	Regret	Requirement on B	Phase I length
m -dim Linear	Theorem 3, Agrawal and Devanur (2016)	$\tilde{O}((\frac{TOPT}{B} + 1)m\sqrt{T})$	$\tilde{\Omega}(mT^{3/4})$	$O(m\sqrt{T})$
	Corollary 1, Sivakumar et al. (2022)		$\Omega(m^{2/3}T^{3/4})$	$\tilde{O}(m^{2/3}\sqrt{T})$
	Theorem 4.1	$\tilde{O}((\frac{TOPT}{B} + 1)\sqrt{KdmT})$	$\tilde{\Omega}((md)^{1/3}(KT)^{3/4})$	$\tilde{O}((md)^{1/3}\sqrt{KT})$
p -Nonparametric	Theorem 4.1	$\tilde{O}((\frac{TOPT}{B} + 1)\sqrt{d}(KT)^{\frac{1+p}{2+p}})$	$\tilde{\Omega}(d^{\frac{2+p}{6+2p}}(KT)^{\frac{3+p}{4+2p}})$	$\tilde{O}(d^{\frac{2+p}{6+2p}}(KT)^{\frac{1+p}{2+p}})$

$$\mathcal{M}(T_0) := \sqrt{K(\mathcal{E}_{T_0}(\mathcal{F}) + d\mathcal{E}_{T_0}(\mathcal{G})) + 4\frac{\log(Td)}{T_0}},$$

and $\mathcal{E}_{T_0}(\mathcal{F}) := \mathcal{E}_{1/T}(T_0; \mathcal{F})$, $\mathcal{E}_{T_0}(\mathcal{G}) := \mathcal{E}_{1/T}(T_0; \mathcal{G})$.

The above linear programming can be seen as an empirical approximation to the static programming (1). We further present a lemma that gives an estimator of $TOPT/B$ whose error is bounded by the estimation error of $\mathcal{R}_{\text{est}}^r$ and $\mathcal{R}_{\text{est}}^c$:

Lemma 4.1. *Denoting the optimal value of (6) by $\widehat{OPT}(T_0)$ and set $Z = \frac{T}{B}(\widehat{OPT}(T_0) + \mathcal{M}(T_0))$, we have with probability at least $1 - O(1/T^2)$,*

$$\frac{TOPT}{B} \leq Z \leq \left(\frac{6T\mathcal{M}(T_0)}{B} + 1\right)\left(\frac{TOPT}{B} + 1\right)$$

In particular, $Z \lesssim \frac{TOPT}{B}$ if $B = \Omega(T(\mathcal{M}(T_0)))$.

With the estimation error guarantee of Z , we can obtain the following regret guarantee of Algorithm 2 by a modification of the proof of Theorem 3.1:

Theorem 4.1. *When $B > \max\{(K + 2)T_0, T\mathcal{M}(T_0)\}$, under Assumption 2.1 and Assumption 4.1, Algorithm 2 achieves the regret*

$$\begin{aligned} \text{Reg}(T) \lesssim & \left(\frac{TOPT}{B} + 1\right)\sqrt{KT[\text{Reg}_{sq}^r(T) + \text{Reg}_{sq}^c(T) + 1]} \\ & + \left(\frac{TOPT}{B} + 1\right)KT_0 \end{aligned}$$

The regret bound in Theorem 4.1 reveals the lower bound requirement on B and provides a guideline on selecting the exploration length T_0 for general \mathcal{F}, \mathcal{G} . We summarize results of T_0 selection and requirements of B when \mathcal{F} and $\tilde{\mathcal{G}}$ are linear classes and nonparametric classes in Table 2 and leave the proof to Appendix. While Theorem 3.1 can apply to more general \mathcal{F} and \mathcal{G} as in Table 1, we present the result with $\mathcal{F} = \mathcal{G}$ for simplicity. In particular, in the linear CBwK setting, both the regret bound and the requirement on B have a better dependency on m than previous results. On the other hand, compared with previous results in the linear setting, our algorithm has an additional dependency on K in the Phase I length. The reason for such dependency is that we take a uniform exploration for general \mathcal{F} and \mathcal{G} instead of an adaptive exploration procedure as in the linear setting (Agrawal and Devanur, 2016). Developing more efficient algorithms of estimating \widehat{OPT} for general function classes is an interesting future direction.

5 CONCLUSION

In this paper, we present a new algorithm for CBwK problem with general reward and cost classes. Our algorithm provides a reduction from CBwK problems to online regression problems. By providing the regret upper bound that matches the lower bound, we demonstrate the optimality of our algorithm for various function classes.

There are several future directions that can be explored. First, our assumption on the cost regression oracle will implicitly lead to an extra \sqrt{d} factor in many examples. One future direction is to relax this assumption to improve the dependency on d . Another related open question is whether a similar reduction from CBwK problems to offline regression problems is possible as in the CB setting (Simchi-Levi and Xu, 2021). Finally, extensions of our framework to the misspecified setting (Foster et al., 2020) and large action space settings (Zhu et al., 2022) are also promising directions to explore.

Acknowledgements

The authors would like to thank Xiaocong Xu for helpful discussions and reviewers for valuable suggestions. This work was supported by HKUST IEG19SC04, the Project of Hetao Shenzhen-HKUST Innovation Cooperation Zone HZQB-KCZYB-2020083, the Guangdong-Hong Kong-Macao Joint Laboratory for Data-Driven Fluid Dynamics, and Hong Kong Research Grant Council (HKRGC) Grant 16214121, 16208120.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24.
- Abe, N. and Long, P. M. (1999). Associative reinforcement learning using linear probabilistic concepts. In *ICML*, pages 3–11. Citeseer.
- Agarwal, A., Hsu, D., Kale, S., Langford, J., Li, L., and Schapire, R. (2014). Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pages 1638–1646. PMLR.

- Agrawal, S. and Devanur, N. (2016). Linear contextual bandits with knapsacks. *Advances in Neural Information Processing Systems*, 29.
- Agrawal, S. and Devanur, N. R. (2014a). Bandits with concave rewards and convex knapsacks. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 989–1006.
- Agrawal, S. and Devanur, N. R. (2014b). Fast algorithms for online stochastic convex programming. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1405–1424. SIAM.
- Agrawal, S., Devanur, N. R., and Li, L. (2016). An efficient algorithm for contextual bandits with knapsacks, and an extension to concave objectives. In *Conference on Learning Theory*, pages 4–18. PMLR.
- Ai, R., Chen, Z., Deng, X., Pan, Y., Wang, C., and Yang, M. (2022). On the re-solving heuristic for (binary) contextual bandits with knapsacks. *arXiv preprint arXiv:2211.13952*.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256.
- Badanidiyuru, A., Kleinberg, R., and Slivkins, A. (2018). Bandits with knapsacks. *Journal of the ACM (JACM)*, 65(3):1–55.
- Badanidiyuru, A., Langford, J., and Slivkins, A. (2014). Resourceful contextual bandits. In *Conference on Learning Theory*, pages 1109–1134. PMLR.
- Balseiro, S. R., Lu, H., and Mirrokni, V. (2022). The best of many worlds: Dual mirror descent for online allocation problems. *Operations Research*.
- Bietti, A., Agarwal, A., and Langford, J. (2021). A contextual bandit bake-off. *J. Mach. Learn. Res.*, 22:133–1.
- Castiglioni, M., Celli, A., and Kroer, C. (2022). Online learning with knapsacks: the best of both worlds. *arXiv preprint arXiv:2202.13710*.
- Chen, X., Lyu, J., Wang, Y., and Zhou, Y. (2022). Fairness-aware network revenue management with demand learning. *arXiv preprint arXiv:2207.11159*.
- Chu, W., Li, L., Reyzin, L., and Schapire, R. (2011). Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings.
- Dudik, M., Hsu, D., Kale, S., Karampatziakis, N., Langford, J., Reyzin, L., and Zhang, T. (2011). Efficient optimal learning for contextual bandits. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 169–178.
- Elmachtoub, A. N. and Grigas, P. (2022). Smart “predict, then optimize”. *Management Science*, 68(1):9–26.
- Ferreira, K. J., Simchi-Levi, D., and Wang, H. (2018). Online network revenue management using thompson sampling. *Operations research*, 66(6):1586–1602.
- Foster, D., Agarwal, A., Dudik, M., Luo, H., and Schapire, R. (2018). Practical contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 1539–1548. PMLR.
- Foster, D. and Rakhlin, A. (2020). Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR.
- Foster, D. J., Gentile, C., Mohri, M., and Zimmert, J. (2020). Adapting to misspecification in contextual bandits. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11478–11489. Curran Associates, Inc.
- Hazan, E. et al. (2016). Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325.
- Hu, Y., Kallus, N., and Mao, X. (2020). Smooth contextual bandits: Bridging the parametric and non-differentiable regret regimes. In *Conference on Learning Theory*, pages 2007–2010. PMLR.
- Jenatton, R., Huang, J., and Archambeau, C. (2016). Adaptive algorithms for online convex optimization with long-term constraints. In *International Conference on Machine Learning*, pages 402–411. PMLR.
- Krishnamurthy, A., Agarwal, A., and Dudik, M. (2016). Contextual semibandits via supervised learning oracles. *Advances In Neural Information Processing Systems*, 29.
- Langford, J. and Zhang, T. (2007). The epoch-greedy algorithm for contextual multi-armed bandits. *Advances in neural information processing systems*, 20(1):96–1.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.
- Liu, H. and Grigas, P. (2022). Online contextual decision-making with a smart predict-then-optimize method. *arXiv preprint arXiv:2206.07316*.
- Mahdavi, M., Jin, R., and Yang, T. (2012). Trading regret for efficiency: online convex optimization with long term constraints. *The Journal of Machine Learning Research*, 13(1):2503–2528.
- Miao, S. and Wang, Y. (2021). Network revenue management with nonparametric demand learning: \sqrt{T} -regret and polynomial dimension dependency. Available at SSRN 3948140.
- Rakhlin, A., Shamir, O., and Sridharan, K. (2012). Making gradient descent optimal for strongly convex stochastic

- optimization. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1571–1578.
- Rakhlin, A., Sridharan, K., and Tsybakov, A. B. (2017). Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 23(2):789–824.
- Rigollet, P. and Zeevi, A. (2010). Nonparametric bandits with covariates. *arXiv preprint arXiv:1003.1630*.
- Shalev-Shwartz, S. et al. (2012). Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194.
- Simchi-Levi, D. and Xu, Y. (2021). Bypassing the monster: A faster and simpler optimal algorithm for contextual bandits under realizability. *Mathematics of Operations Research*.
- Sivakumar, V., Zuo, S., and Banerjee, A. (2022). Smoothed adversarial linear contextual bandits with knapsacks. In *International Conference on Machine Learning*, pages 20253–20277. PMLR.
- Slivkins, A. (2011). Contextual bandits with similarity information. In *Proceedings of the 24th annual Conference On Learning Theory*, pages 679–702. JMLR Workshop and Conference Proceedings.
- Slivkins, A. et al. (2019). Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286.
- Slivkins, A. and Foster, D. (2022). Efficient contextual bandits with knapsacks via regression. *arXiv preprint arXiv:2211.07484*.
- Zhu, Y., Foster, D. J., Langford, J., and Mineiro, P. (2022). Contextual bandits with large action spaces: Made practical. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S., editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 27428–27453. PMLR.

A PROOF OF RESULTS IN SECTION 3

A.1 Proof of Theorem 3.1

We would prove the following more general form of Theorem 3.1

Theorem A.1. *Under Assumption 2.1, suppose $Z \geq TOPT/B$ and the dual-update step is running over Λ with radius Z , if the output of \mathcal{R}^r and \mathcal{R}^c satisfy (2) and (3), respectively, denote*

$$\gamma = \sqrt{KT / (\text{Reg}_{S_q}^r(T) + Z^2 \text{Reg}_{S_q}^c(T) + 4 \log(2T))},$$

then SquareCBwK achieves the regret

$$\text{Reg}(T) \lesssim (Z + 1) \sqrt{KT \cdot (\text{Reg}_{S_q}^r(T) + \text{Reg}_{S_q}^c(T) + \log(dT))}.$$

In particular, since it always hold that $TOPT/B \leq T/B$, Theorem 3.1 is a special case of Theorem A.1 with setting $Z = T/B$.

Proof of Theorem A.1. Denote $\ell_t(a) := f^*(x_t, a) + \lambda_t^T (B/T \cdot \mathbf{1} - \mathbf{g}^*(x_t, a))$ and $p^*(x)$ the optimal solution of static programming (1), $p_{t,a}^* := p_a^*(x_t)$, then we have

$$\begin{aligned} & \sum_{t=1}^{\tau} \left[\sum_{a'=1}^K p_{t,a'}^* \ell_t(a') - \sum_{a=1}^K p_{t,a} \ell_t(a) \right] \\ &= \sum_{t=1}^{\tau} \sum_{a'=1}^K p_{t,a'}^* [\ell_t(a') - \sum_{a=1}^K p_{t,a} \ell_t(a)] \\ &= \sum_{t=1}^{\tau} \sum_{a'=1}^K p_{t,a'}^* \left(\sum_{a=1}^K p_{t,a} [\ell_t(a') - \ell_t(a) - \frac{\gamma}{4} (\hat{\ell}_t(a) - \ell_t(a))^2 + \frac{\gamma}{4} (\hat{\ell}_t(a) - \ell_t(a))^2] \right), \end{aligned}$$

where we also denote $\hat{\ell}_t(a) := \hat{\ell}_{t,a} = \hat{r}_{t,a} + \lambda_t^T (B/T \cdot \mathbf{1} - \hat{\mathbf{c}}_{t,a})$. Notice that denote $a_t^* := \arg\max_{a \in [K]} \ell_t(a)$, we have

$$\sum_{a=1}^K p_{t,a} [\ell_t(a') - \ell_t(a) - \frac{\gamma}{4} (\hat{\ell}_t(a) - \ell_t(a))^2] \leq \sum_{a=1}^K p_{t,a} [\ell_t(a_t^*) - \ell_t(a) - \frac{\gamma}{4} (\hat{\ell}_t(a) - \ell_t(a))^2] \leq \frac{2K}{\gamma}.$$

Where in the second inequality we used Lemma 3 in Foster and Rakhlin (2020). Thus we get

$$\sum_{t=1}^{\tau} \left[\sum_{a'=1}^K p_{t,a'}^* \ell_t(a') - \sum_{a=1}^K p_{t,a} \ell_t(a) \right] \leq \sum_{t=1}^{\tau} \left(\frac{2K}{\gamma} + \sum_{a=1}^K \frac{\gamma}{4} p_{t,a} (\hat{\ell}_t(a) - \ell_t(a))^2 \right).$$

Now we would control $\sum_{t=1}^{\tau} \sum_{a=1}^K \frac{\gamma}{4} p_{t,a} (\hat{\ell}_t(a) - \ell_t(a))^2$: To deal with the stopping time τ , we establish a concentration result of $\sum_{s=1}^t \sum_{a=1}^K \frac{\gamma}{4} p_{s,a} (\hat{\ell}_s(a) - \ell_s(a))^2$ uniformly for all $1 \leq t_0 \leq T$ via the following variant of Freedman's inequality:

Lemma A.1 (Rakhlin et al. (2012), Lemma 3). *Let $\{D_s\}_{s=1}^T$ be a martingale difference sequence w.r.t. filtration $\{\mathcal{F}_s\}_{s=1}^T$ and with a uniform upper bound b . Let V denote the sum of conditional variances,*

$$V_s = \sum_{i=1}^s \text{Var}(D_i | \mathcal{F}_{i-1}),$$

then for any $\delta < 1/e$ and $T \geq 4$,

$$\mathbb{P} \left(\left| \sum_{s=1}^t D_s \right| > 2 \max \{ 2\sqrt{V_s}, b\sqrt{\log(1/\delta)} \} \sqrt{\log(1/\delta)} \text{ for some } t \leq T \right) \leq 2\delta \log(T).$$

Now noticing that for $\mathcal{F}_{t-1} := \sigma(\{x_t\} \cup \{x_s, a_s, r_s, \mathbf{c}_i\}_{s \leq t-1})$, denote $M_s := (\hat{\ell}_s(a_s) - \ell_s(a_s))^2$, we have

$$D_s := M_s - \mathbb{E}[M_s | \mathcal{F}_{s-1}] = (\hat{\ell}_s(a_s) - \ell_s(a_s))^2 - \sum_{a=1}^K p_{s,a} (\hat{\ell}_s(a) - \ell_s(a))^2, 1 \leq s \leq T$$

is a martingale difference sequence uniformly bounded by $4Z^2 + 4$ with respect to \mathcal{F}_s , i.e. $|D_s| \leq 4Z^2 + 4$, $\{D_s\}_{s=1}^T$ is adaptive to $\{\mathcal{F}_s\}_{s=1}^T$ and $\mathbb{E}[D_s | \mathcal{F}_{s-1}] = 0$. By

$$\text{Var}(D_i | \mathcal{F}_{i-1}) \leq \mathbb{E}[D_i^2 | \mathcal{F}_{i-1}] \leq (4Z^2 + 4) \mathbb{E}[|D_i| | \mathcal{F}_{i-1}] \leq (8Z^2 + 8) \mathbb{E}[M_i | \mathcal{F}_{i-1}].$$

We have by Lemma A.1, with probability at least $1 - O(1/T^2)$,

$$\begin{aligned} \sum_{s=1}^t \sum_{a=1}^K p_{s,a} (\hat{\ell}_s(a) - \ell_s(a))^2 &\leq 8\sqrt{\log T} \max \left\{ \sqrt{(8Z^2 + 8) \sum_{i=1}^s \sum_{a=1}^K p_{s,a} (\hat{\ell}_s(a) - \ell_s(a))^2}, (4Z^2 + 4) \sqrt{\log T} \right\} \\ &\quad + \sum_{s=1}^t (\hat{\ell}_s(a_s) - \ell_s(a_s))^2 \end{aligned}$$

holds for uniformly for $1 \leq t \leq T$.

When $\sum_{s=1}^t \sum_{a=1}^K p_{s,a} (\hat{\ell}_s(a) - \ell_s(a))^2 > (2Z^2 + 2) \log T$, we have by

$$A \leq B + C\sqrt{A} \implies (\sqrt{A} - \frac{C}{2})^2 \leq B + \frac{C^2}{4} \implies A \leq 2B + C^2, \quad \forall A, B, C > 0,$$

where

$$A = \sum_{s=1}^t \sum_{a=1}^K p_{s,a} (\hat{\ell}_s(a) - \ell_s(a))^2, \quad B = \sum_{s=1}^t (\hat{\ell}_s(a_s) - \ell_s(a_s))^2, \quad C = 8\sqrt{(8Z^2 + 8) \log T}$$

then implies with probability at least $1 - O(1/T^2)$,

$$\begin{aligned} \sum_{s=1}^t \sum_{a=1}^K p_{s,a} (\hat{\ell}_s(a) - \ell_s(a))^2 &\leq \max\{(2Z^2 + 2) \log T, 2 \sum_{s=1}^t (\hat{\ell}_s(a_s) - \ell_s(a_s))^2 + (512Z^2 + 512) \log T\} \\ &= 2 \sum_{s=1}^t (\hat{\ell}_s(a_s) - \ell_s(a_s))^2 + (512Z^2 + 512) \log T, \quad \forall 1 \leq t \leq T. \end{aligned}$$

Finally, by our assumption on $\mathcal{R}^r, \mathcal{R}^c$,

$$\begin{aligned} \sum_{s=1}^t (\hat{\ell}_s(a_s) - \ell_s(a_s))^2 &= \sum_{s=1}^t (\hat{r}_{s,a_s} - f^*(x_s, a_s) + \boldsymbol{\lambda}_s^T (\mathbf{g}^*(x_s, a_s) - \hat{\mathbf{c}}_{s,a_s}))^2 \\ &\leq \sum_{s=1}^t 2 \left([\hat{r}_{s,a_s} - f^*(x_s, a_s)]^2 + Z^2 \|\mathbf{g}^*(x_s, a_s) - \hat{\mathbf{c}}_{s,a_s}\|_\infty^2 \right). \\ &\leq 2\text{Reg}_{\text{sq}}^r(T) + 2Z^2 \text{Reg}_{\text{sq}}^c(T). \end{aligned}$$

we get with probability at least $1 - O(1/T^2)$,

$$\sum_{t=1}^\tau \sum_{a=1}^K p_{t,a} \ell_t(a) \geq \sum_{t=1}^\tau \sum_{a=1}^K p_{t,a}^* \ell_t(a) - O\left(\frac{2KT}{\gamma} + \frac{\gamma}{4} [\text{Reg}_{\text{sq}}^r(T) + Z^2 \text{Reg}_{\text{sq}}^c(T) + (Z^2 + 1) \log T]\right).$$

Applying Azuma-Hoeffding inequality to $\sum_{t=1}^T \ell_t(a_t)$, and by our selection of γ , we get with probability at least $1 - O(1/T^2)$,

$$\sum_{t=1}^\tau \ell_t(a_t) \geq \sum_{t=1}^\tau \sum_{a=1}^K p_{t,a}^* \ell_t(a) - O((Z+1) \sqrt{KT [\text{Reg}_{\text{sq}}^r(T) + \text{Reg}_{\text{sq}}^c(T)]})$$

$$\begin{aligned}
 &= \sum_{t=1}^{\tau} \left[\left\langle \sum_{a=1}^K p_{t,a}^* f^*(x_t, a) + \lambda_t^T \left(\frac{B}{T} \cdot \mathbf{1} - \sum_{a=1}^K p_{t,a}^* \mathbf{g}^*(x_t, a) \right) \right\rangle \right. \\
 &\quad \left. - O((Z+1)\sqrt{KT[\text{Reg}_{sq}^r(T) + \text{Reg}_{sq}^c(T) + \log T]}) \right].
 \end{aligned}$$

Now consider a new filtration $\mathcal{F}'_{t-1} = \sigma(\{x_s, a_s, r_s, \mathbf{c}_s\}_{s=1}^{t-1})$, we have denote

$$D'_t := \sum_{a=1}^K p_{t,a}^* f^*(x_t, a) + \lambda_t^T \left(\frac{B}{T} \cdot \mathbf{1} - \sum_{a=1}^K p_{t,a}^* \mathbf{g}^*(x_t, a) \right) - \text{OPT},$$

then D'_t is adaptive to \mathcal{F}'_t , $|D'_t| \leq Z + 2$, and $\mathbb{E}[D'_t | \mathcal{F}'_{t-1}] \geq 0$. By the Azuma-Hoeffding inequality, we get with probability at least $1 - O(1/T^2)$,

$$\sum_{s=1}^t D'_s \gtrsim -(Z+1)\sqrt{T \log T}, \quad \forall 1 \leq t \leq T.$$

That leads to with probability at least $1 - O(1/T^2)$,

$$\sum_{t=1}^{\tau} \ell_t(a_t) \geq \tau \text{OPT} - O((Z+1)\sqrt{KT[\text{Reg}_{sq}^r(T) + \text{Reg}_{sq}^c(T) + \log T]}).$$

On the other hand, by Lemma 2.2, for any fixed $\lambda \in \Lambda$ we have,

$$\sum_{t=1}^{\tau} f^*(x_t, a_t) + \lambda_t^T (B/T \cdot \mathbf{1} - \mathbf{c}_{t,a_t}) \leq \sum_{t=1}^{\tau} f^*(x_t, a_t) + \lambda^T (B/T \cdot \mathbf{1} - \mathbf{c}_{t,a_t}) + O(Z\sqrt{T \log d}).$$

By applying Azuma-Hoeffding inequality to summation of $\lambda_t^T (B/T \cdot \mathbf{1} - \mathbf{c}_{t,a_t})$ with respect to $\tilde{\mathcal{F}}_{t-1} := \sigma\{\mathcal{F}_t \cup \{a_t\}\}$, we get with probability at least $1 - O(1/T^2)$,

$$\sum_{s=1}^t \lambda_s^T (B/T \cdot \mathbf{1} - \mathbf{c}_{s,a_s}) \geq \sum_{s=1}^t \lambda_s^T (B/T \cdot \mathbf{1} - \mathbf{g}^*(x_s, a_s)) - O((Z+1)\sqrt{T \log(T)}), \quad \forall 1 \leq t \leq T.$$

Combining all results together, we get with probability at least $1 - O(1/T^2)$,

$$\sum_{t=1}^{\tau} f^*(x_t, a_t) + \lambda^T (B/T \cdot \mathbf{1} - \mathbf{c}_{t,a_t}) \geq \tau \cdot \text{OPT} - O((Z+1)\sqrt{KT[\text{Reg}_{sq}^r(T) + \text{Reg}_{sq}^c(T) + \log(dT)]}), \quad \forall \lambda \in \Lambda.$$

Now

Case1: if $\tau = T$, we get the desired regret bound by selecting $\lambda = \mathbf{0}$.

Case2: If $\tau < T$, there exists some resource j running out, i.e. $\sum_{t=1}^{\tau} (c_{t,a_t})_j > B - 1$, then letting $\lambda = Z\mathbf{e}_j$, we get

$$\begin{aligned}
 \sum_{t=1}^{\tau} [f^*(x_t, a_t) + \lambda^T (B/T \cdot \mathbf{1} - \mathbf{c}_{t,a_t})] &< \sum_{t=1}^{\tau} f^*(x_t, a_t) + Z(\tau B/T - B + 1) \\
 &\leq \sum_{t=1}^{\tau} f^*(x_t, a_t) + \text{OPT}(\tau - T) + Z.
 \end{aligned}$$

where the second inequality is by $\tau/T < 1$ and $Z \geq T\text{OPT}/B$. This inequality leads to with probability at least $1 - O(1/T^2)$

$$\sum_{t=1}^{\tau} f^*(x_t, a_t) > T\text{OPT} - Z - O((Z+1)\sqrt{KT[\text{Reg}_{sq}^r(T) + \text{Reg}_{sq}^c(T) + \log(dT)]}).$$

Denote the event that above inequality holds as \mathcal{Q} , then

$$\text{Reg}(T) \leq \mathbb{E}[(T\text{OPT} - \sum_{t=1}^{\tau} f^*(x_t, a_t))(\mathbf{1}_{\mathcal{Q}} + \mathbf{1}_{\mathcal{Q}^c})]$$

$$\begin{aligned} &\leq \mathbb{E}[(TOPT - \sum_{t=1}^{\tau} f^*(x_t, a_t)) \mathbf{1}_{\mathcal{Q}}] + T\mathbb{P}(Q^c) \\ &\lesssim (Z+1)\sqrt{KT[\text{Reg}_{sq}^r(T) + \text{Reg}_{sq}^c(T) + \log(dT)]}. \end{aligned}$$

Thus the claim holds. \square

A.2 Proof of Theorem 3.2

Proof. When $\text{Reg}_{\text{CB}}(T; \mathcal{F}) > \alpha T$, one can set $B = T$. Then the problem is exactly the unconstrained CB problem with reward class \mathcal{F} , whose regret is lower bounded by $\text{Reg}_{\text{CB}}(T; \mathcal{F})$.

Now we would focus on the case $\alpha T > \text{Reg}_{\text{CB}}(T; \mathcal{F})$: Firstly assume W.L.O.G. αT is not $O(1)$, otherwise the lower bound result is trivial.

For any fixed π , consider the following instance of CBwK:

1. At every round t , x_t is sampled i.i.d. from P_0 .
2. By condition of Theorem 3.2, there exists some $f_0 \in \mathcal{F}$ satisfying $f_0(x, 0) = f_0(x, 1) \geq c_0$ a.s. P_0 . We set the reward as $r_{t,a} = f_0(x_t, a)$.
3. By \mathcal{G} is α -separable, there exists some $g_0 \in \mathcal{G}$ and distribution Q_0 of costs such that

$$p := \mathbb{E}_{P_0}[\min_a g_0(x, a)] \geq \frac{1}{4}, \quad \mathbb{E}_{P_0}[\max_a g_0(x, a)] \geq p + \alpha.$$

and there exists some $c_1 > 0$ s.t.

$$\mathbb{E}_{P_0, Q_0}[\sum_{t=1}^{T/8} \mathbf{1}\{\pi_t(x_t) \neq \text{argmin}_a g_0(x, a)\}] > c_1 T.$$

we let the cost be generated from Q_0 .

For such instance, we have the regret of π is lower bounded by $c_0(\mathbb{E}[\tau^*] - \mathbb{E}[\tau^\pi])$, with τ^π the stopping time of π and τ^* the stopping time of π^* , with

$$\pi_t^*(x_t) = \text{argmax}_a g_0(x_t, a).$$

Now we would bound $\mathbb{E}[\tau^* - \tau^\pi]$ from below:

Lower bound of $\mathbb{E}[\tau^*]$: For U_i i.i.d. from the distribution of $c_{1, \pi^*(1, x_1)}$, we have then τ^* follows the same distribution as $\tau_U \wedge T$. Consider the first-hitting times

$$\inf\{t \geq 0 : \sum_{s=1}^t U_s > B - 1\}, \quad \tau_V := \inf\{t \geq 0 : \sum_{s=1}^t V_s > B - 1\},$$

notice that by $\mathbb{E}[U_i] = p \geq 1/4$ and $U_i \in [0, 1]$, we have $\tau_U < \infty$ almost surely. Thus by Wald's equation,

$$B - 1 < \mathbb{E}[\sum_{t=1}^{\tau} U_t] = \mathbb{E}[U_1]\mathbb{E}[\tau_U],$$

that leads to $\mathbb{E}[\tau_U] \geq \frac{B-1}{p}$. Notice that for $\tau'_U = \tau_U \wedge T$, we have

$$\begin{aligned} \mathbb{E}[\tau_U] - \mathbb{E}[\tau'_U] &= \sum_{t \geq 0} \mathbb{P}(\tau_U > 0) - \sum_{t \geq 0} \mathbb{P}(\tau'_U > 0) \\ &= \sum_{t \geq T} \mathbb{P}(\tau_U > t) \end{aligned}$$

$$\begin{aligned} &\leq \sum_{t \geq T} \mathbb{P}\left(\sum_{s=1}^m U_s < B-1, \forall m \leq t\right) \\ &\leq \sum_{t \geq T} \mathbb{P}\left(\sum_{s=1}^t U_s < B-1\right). \end{aligned}$$

Notice that by Hoeffding's inequality and $p \geq 1/4$,

$$\mathbb{P}\left(\sum_{s=1}^t U_s < \frac{t}{4} - u\right) \leq \mathbb{P}\left(\left|\sum_{s=1}^t U_s - pt\right| > u\right) \lesssim \exp(-Cu^2/t),$$

thus by $B = T/8$,

$$\mathbb{P}\left(\sum_{s=1}^t U_s < B-1\right) \leq \mathbb{P}\left(\sum_{s=1}^t U_s < \frac{t}{4} - \frac{t}{8}\right) \lesssim \exp(-ct), \quad \forall t > T.$$

That leads to

$$\mathbb{E}[\tau_U] - \mathbb{E}[\tau'_U] \lesssim \sum_{t \geq T} \exp(-ct) \lesssim \exp(-cT).$$

Thus we get

$$\mathbb{E}[\tau^*] = \mathbb{E}[\tau'_U] \geq \frac{B-1}{p} - O(\exp(-cT)).$$

Upper bound of $\mathbb{E}[\tau^\pi]$: Recall the notation $\mathcal{F}_{t-1} := \sigma(\{x_t\} \cup \{x_s, a_s, r_s, c_s\}_{s=1}^{t-1})$, since every pulling with $\pi_t(x_t) \neq \operatorname{argmin}_a g_0(x_t, a)$ will incur a cost at least $p + \alpha$ in expectation, we have

$$\mathbb{E}[c_{t, \pi(x_t)} | \mathcal{F}_{t-1}] \geq p + \alpha - \alpha \mathbb{E}[\mathbf{1}\{\pi(x_t) = \pi_{\epsilon_0}^*(x_t)\} | \mathcal{F}_{t-1}].$$

Thus if we denote

$$\tilde{D}_t := c_{t, \pi(x_t)} - (p + \alpha) + \alpha \mathbf{1}\{\pi(x_t) = \pi^*(x_t)\}, t \geq 1, \quad \tilde{D}_0 = 0.$$

Then \tilde{D}_t is a sub-martingale difference sequence with respect to \mathcal{F}_t , i.e. $\tilde{M}_t := \sum_{s=0}^t \tilde{D}_s$ is a sub-martingale with respect to \mathcal{F}_t . Denote $\tilde{\tau}' := \inf\{t \geq 0, \sum_{t=0}^t c_{t, \pi(x_t)} > B-1\}$, then $\tilde{\tau}' < \infty$ a.s., thus we have by optional stopping theorem, $\mathbb{E}[\tilde{M}_{\tilde{\tau}'}] \geq 0$.

If we denote C_t the total cost of π up to time t and $N(t)$ the total times of pulling on or before t so that $\pi_t \neq \pi^*(x_t)$, then

$$\mathbb{E}[\tilde{M}_{\tilde{\tau}'}] = \mathbb{E}[C_{\tilde{\tau}'} - (p + \alpha)\tilde{\tau}' + \alpha N(\tilde{\tau}')] \geq 0, \quad (7)$$

thus by $\tilde{\tau}' \geq T/8$

$$B \geq \mathbb{E}[(p + \alpha)\tilde{\tau}' - \alpha N(T/8)] \geq (p + \alpha)\mathbb{E}[\tilde{\tau}'] - c_1 \alpha T,$$

i.e. $\mathbb{E}[\tilde{\tau}'] \leq \frac{B}{p + \alpha} - \frac{c_1 \alpha}{p + \alpha} T$.

Now since $\tau^\pi = \tilde{\tau}' \wedge T$, we have

$$\mathbb{E}[\tau^\pi] \leq \mathbb{E}[\tilde{\tau}'] \leq \frac{B}{p + \alpha} - \frac{c_1 \alpha}{p + \alpha} T.$$

Lower bound of $\mathbb{E}[\tau^* - \tau^\pi]$ Combing bounds for $\mathbb{E}[\tau^\pi], \mathbb{E}[\tau^*]$ together, we have

$$\begin{aligned} \mathbb{E}[\tau^* - \tau^\pi] &\geq \left(\frac{1-1/B}{p} - \frac{1}{p+\alpha}\right)B + \frac{c_1 \alpha T}{p+\alpha} - O(\exp(-cT)) \\ &\gtrsim \alpha T + (p + \alpha - \frac{p + \alpha}{B} - p)B - O(\exp(-cT)) \\ &\gtrsim \alpha T - O(1). \end{aligned}$$

Noticing that αT dominates the last term, thus the claim holds. \square

A.3 Detail of Results in Section 3.3

A.3.1 Construction of \mathcal{R}^c from $\tilde{\mathcal{R}}^c$

By $\mathcal{G} = \tilde{\mathcal{G}}^d$, we can denote the underlying expected cost function $\mathbf{g}^*(x, a) = (g_1^*(x, a), \dots, g_d^*(x, a))$. By our assumption on $\tilde{\mathcal{R}}^c$, we have running $\tilde{\mathcal{R}}^c$ over a online regression problem with underlying function $g_i^*(\cdot, \cdot)$ will generate a sequence of predictors $\{\hat{g}_{t,i}(\cdot, \cdot)\}_{t=1}^T$ so that

$$\sum_{t=1}^T (\hat{g}_{t,i}(x_t, a_t) - g_i^*(x_t, a_t))^2 \leq \widetilde{\text{Reg}}_{\text{sq}}^c(T).$$

So if we set \mathcal{R}^c as the oracle with output $\hat{\mathbf{g}}_t(x, a) := (\hat{g}_{t,1}(x, a), \dots, \hat{g}_{t,d}(x, a))$ at each round t , then it satisfies that

$$\sum_{t=1}^T \|\hat{\mathbf{g}}_t(x_t, a_t) - \mathbf{g}_t^*(x_t, a_t)\|_\infty^2 \leq \sum_{i=1}^d \sum_{t=1}^T (\hat{g}_{t,i}(x_t, a_t) - g_i^*(x_t, a_t))^2 \leq d \widetilde{\text{Reg}}_{\text{sq}}^c(T).$$

A.3.2 Generalized Linear CBwK

Selection of Oracles and Upper Regret Bound For the m -dimensional generalized linear class, by Proposition 3.3 of Foster and Rakhlin (2020), the Newtonized GLMtron oracle achieves the $O(m \log T)$ online regression regret. So selecting $\mathcal{R}^r, \tilde{\mathcal{R}}^c$ as GLMtron oracle implies

$$\text{Reg}_{\text{sq}}^r(T) \lesssim m_1 \log T, \quad \text{Reg}_{\text{sq}}^c(T) \lesssim m_2 d \log T.$$

Bringing this result to Theorem 3.1 leads to the desired regret upper bound of generalized linear CBwK.

Regret Lower Bound Result Since Generalized linear CBwK includes linear CBwK as a special case, we focus on proving the lower bound result in linear setting. Formally, we show the following lower bound result for linear CBwK:

Theorem A.2. *Let \mathcal{F}, \mathcal{G} be m_1 -dimensional and m_2 -dimensional linear function classes respectively. Then there exists the selection of feature maps $\phi_1, \phi_2, T > 0, B \geq T/4$, so that for a CBwK instance with time horizon T , budget B and the reward and cost function classes \mathcal{F}, \mathcal{G} with $K = 2, d = 1$, any policy π must have*

$$\mathbb{E}[\text{Reg}_\pi(T)] = \tilde{\Omega}(\sqrt{(m_1 + m_2)T})$$

Proof of Theorem A.2. We would prove Theorem A.3 by verifying the conditions in Theorem 3.2. Our analysis includes three steps:

1. Specify \mathcal{X} and feature maps ϕ_1, ϕ_2 .
2. Verify \mathcal{G} is α -inseparable.
3. Verify the existence of $f_0 \in \mathcal{F}$ which satisfies the condition of Theorem 3.2.
4. Plug the αT and $\text{Reg}_{\text{CB}}(T; \mathcal{F})$ into Theorem 3.2 to get the desired lower bound.

Step1: We specify \mathcal{X} as the a subset $[0, 1]^{2m_1+2m_2}$ defined as following:

$$\mathcal{X} = \{x = (\underbrace{x^{(1)}}_{\mathbb{R}^{m_1}}, \underbrace{x^{(2)}}_{\mathbb{R}^{m_1}}, \underbrace{x^{(3)}}_{\mathbb{R}^{m_2}}, \underbrace{x^{(4)}}_{\mathbb{R}^{m_2}}) : \|x^{(i)}\|_2 \leq 1, i = 1, 2, 3, 4\}.$$

The feature maps $\phi_1(x, a), \phi_2(x, a)$ is defined as

$$\phi_1(x, 0) = x^{(1)}, \phi_1(x, 1) = x^{(2)}, \phi_2(x, 0) = x^{(3)}, \phi_2(x, 1) = x^{(4)}.$$

With such construction, determine the distribution over x is equivalent to determine the distribution of $(x^{(1)}, x^{(2)})$ and $(x^{(3)}, x^{(4)})$, we will determine them in Step2 and Step3.

Step 2: As constructed in Step 1, we need only construct the distribution of $(\phi_2(x, 0), \phi_2(x, 1))$. Our construction is motivated by a deterministic construction used in Chu et al. (2011). Consider the following subset of $[0, 1]^{2m_2}$: W.L.O.G. assume $r = (m_2 - 1)/2$ is an integer, let

$$\tilde{\mathcal{X}} := \{\tilde{x}_i := \frac{1}{2}e_1 + \frac{1}{2}e_{2i} + \frac{1}{2}e_{m_2+1} + \frac{1}{2}e_{m_2+2i+1}, i \in [r]\},$$

then when $(x^{(3)}, x^{(4)}) \in \tilde{\mathcal{X}}$, we have

$$\phi_2(x, 0) = (\frac{1}{2}, 0, \dots, \frac{1}{2}, 0, \dots, 0), \phi_2(x, 1) = (\frac{1}{2}, 0, \dots, 0, \frac{1}{2}, \dots, 0).$$

We let the distribution of $(x^{(3)}, x^{(4)}) \tilde{P}_0$ as the uniform distribution over $\tilde{\mathcal{X}}$.

On the other hand, we construct the subset \mathcal{U}_γ of \mathcal{G} as following: For every $\epsilon \in \{-1, 1\}^r$, let

$$(\theta_\epsilon)_1 = \frac{1}{2}, (\theta_\epsilon)_{2i, 2i+1} = \begin{cases} (\gamma, 0) & \text{if } \epsilon_i = -1, \\ (0, \gamma) & \text{if } \epsilon_i = 1, \end{cases} \quad i \in [r], \quad (\theta_\epsilon)_j = 0 \text{ for other } j.$$

we set

$$\mathcal{U}_\gamma := \{(x, a) \rightarrow \langle \phi_2(x, a), \theta_\epsilon \rangle, \quad \epsilon \in \{-1, 1\}^r\},$$

then for any $g \in \mathcal{U}_\gamma$, we have

$$\frac{1}{4} = \mathbb{E}_x[\min_a \phi_2(x, a)] < \mathbb{E}_x[\max_a \phi_2(x, a)] = \frac{1}{4} + \frac{\gamma}{2}$$

If we set $r(x, a)$ generated as Bernoulli($\phi_2(x, a)$) condition on x, a and consider the uniform distribution Q over CB instances with underlying $g^* \in \mathcal{U}_\gamma$, since by our construction $g(\tilde{x}_i, a), g(\tilde{x}_j, a)$ are independently distributed under Q when $\tilde{x}_i \neq \tilde{x}_j$, we have when consider the CB instance with $g \sim Q$, the data generating process can be seen as following, as discussed in Foster and Rakhlin (2020):

1. Sample x_1, \dots, x_T i.i.d. from \tilde{P}_0 , set $S_i := \{t \in [T] : x_t = \tilde{x}_i\}$.
2. For each $i \in [r]$, independently sample a Bernoulli MAB instance \mathcal{P}_i with arm means μ_1, μ_2 such that with probability $\frac{1}{2}$,

$$\mu_1 = \frac{1}{4}, \mu_2 = \frac{1}{4} + \gamma$$

and with probability $\frac{1}{2}$,

$$\mu_2 = \frac{1}{4} + \gamma, \mu_1 = \frac{1}{4}.$$

Now by the same statement as in Foster and Rakhlin (2020), we have selecting $\gamma \propto \sqrt{\frac{r}{T}} \propto \sqrt{\frac{m_2}{T}}$ leads to for any π , there exists some $\epsilon_0 \in \{-1, 1\}^r$ so that for the CB instance with underlying reward function is g_{ϵ_0} ,

$$\frac{\gamma}{2} \mathbb{E}[\sum_{t=1}^T \mathbf{1}\{\operatorname{argmax}_a g_{\epsilon_0}(x, a) \neq \pi_t(x_t)\}] = \mathbb{E}[\sum_{t=1}^T \max_a g_{\epsilon_0}(x_t, a) - g_{\epsilon_0}(x_t, \pi_t(x_t))] \gtrsim \gamma T.$$

Thus \mathcal{G} is $\sqrt{\frac{m_2}{T}}$ -inseparable. Above result also implies a $\Omega(\sqrt{m_2 T})$ lower bound for $\operatorname{Reg}_{\text{CB}}$ over \mathcal{G} .

Step2: We can simply set the distribution \tilde{P}' of $(x^{(1)}, x^{(2)})$ as a constant distribution x_0 with $\frac{1}{4}$ in its first and $m_1 + 1$ -th coordinate and 0 in other coordinates. Then let P_0 be the distribution of $(x^{(1)}, x^{(2)}, x^{(3)}, x^{(4)})$ given by $(x^{(1)}, x^{(2)}) \sim \tilde{P}', (x^{(3)}, x^{(4)}) \sim \tilde{P}$, we have $f_0(x, a) = \langle \phi_1(x, a), \theta_0 \rangle$ with $\theta_0 = x_0$ is a function satisfying the condition of Theorem 3.2.

Step3: By our construction, we have there exists some $f_0 \in \mathcal{F}$ so that $f(x, a) \equiv \frac{1}{4}$ under P_0 , $\operatorname{Reg}_{\text{CB}}(T; \mathcal{F}) = \tilde{\Omega}(\sqrt{m_1 T})$. And \mathcal{G} is $\tilde{\Theta}(\sqrt{\frac{m_2}{T}})$ -inseparable, applying Theorem 3.2 leads to the $\tilde{\Omega}(\max\{\sqrt{m_1 T}, \sqrt{m_2 T}\}) = \tilde{\Omega}(\sqrt{m_1 T} + \sqrt{m_2 T})$ lower bound as desired. \square

A.3.3 Nonparametric CBwK

Selection of Oracles and Upper Regret Bound For the p -nonparametric class, by Theorem 3 of Foster and Rakhlin (2020), a Vovk's aggregation based oracle achieves the $O(KT)^{1-\frac{2}{2+p}}$ online regression regret. So selecting $\mathcal{R}^r, \tilde{\mathcal{R}}^c$ as this oracle implies

$$\text{Reg}_{\text{sq}}^r(T) \lesssim (KT)^{1-\frac{2}{2+p_1}}, \quad \text{Reg}_{\text{sq}}^c(T) \lesssim d(KT)^{1-\frac{2}{2+p_2}}.$$

Bringing this result to Theorem 3.1 leads to the desired regret upper bound of nonparametric CBwK.

Regret Lower Bound Result Formally, we show the following lower bound result for non-parametric CBwK:

Theorem A.3. *Let \mathcal{W}, \mathcal{V} be two function classes consisting of functions from \mathcal{X} to $[0, 1]$, satisfying $H(\mathcal{W}) = \Theta(\varepsilon^{-p_1}), H(\mathcal{V}) = \Theta(\varepsilon^{-p_2})$. Then there exists $T > 0, B \geq T/4$, a slightly modified class $\mathcal{V}', \mathcal{W}'$ with $H(\mathcal{W}') = \tilde{\Theta}(\varepsilon^{-p_1}), H(\mathcal{V}') = \tilde{\Theta}(\varepsilon^{-p_2})$ so that for a CBwK instance with time horizon T , budget B and the reward and cost function classes \mathcal{F}, \mathcal{G} constructed from $\mathcal{V}', \mathcal{W}'$ with $K = 2, d = 1$, any policy π must have*

$$\mathbb{E}[\text{Reg}_\pi] = \tilde{\Omega}(T^{\frac{1+p_1}{2+p_1}} + T^{\frac{1+p_2}{2+p_2}})$$

Theorem A.3 can be seen as a extension of Theorem 2 in Foster and Rakhlin (2020) for non-parametric contextual bandits. One subtlety of these nonparametric results is that, in contrast to the results in linear case, the lower bound is stated over a modification of the original function classes \mathcal{F}, \mathcal{G} . That makes the result slightly different from the classical worst-case lower bound results for CB or CBwK. However, such lower bound implies the optimality of our algorithm with respect to the complexity of the considered function classes, which receive the most attention and often provide tight characterization of the problem in nonparametric setting.

Proof of Theorem A.3. Similarly to the proof of Theorem A.2, we would use Theorem 3.2 to develop Theorem A.3. Our analysis includes three steps:

1. Construct a modification \mathcal{V}' of \mathcal{V} so that its corresponding cost function class \mathcal{G} satisfies the α -inseparable property.
2. Construct a modification \mathcal{W}' of \mathcal{W} so that its corresponding cost function class \mathcal{F} containing some $f_0 \in \mathcal{F}$ that satisfies the condition of Theorem 3.2.
3. Plug the αT and $\text{Reg}_{\text{CB}}(T; \mathcal{F})$ into Theorem 3.2 to get the desired lower bound.

Step1: The construction of \mathcal{G} in Step 1 is from the construction used in the proof of Theorem 2 in Foster and Rakhlin (2020). And we restate several key steps for completeness.

For any $0 < \gamma \leq \frac{1}{4}$ and \mathcal{V} with $H(\mathcal{V}) = \Theta(\varepsilon^{-p_2})$, by the argument in Foster and Rakhlin (2020), we can find $m = \Theta(\gamma^{-p_2})$ distinct $x^{(1)}, \dots, x^{(m)} \in \mathcal{X}$ and a class $\mathcal{V}' \supset \mathcal{V}$ containing $\{h_\epsilon\}_{\epsilon \in \{-1, 1\}^m}$ such that

$$h_\epsilon : \mathcal{X} \rightarrow [0, 1], \quad h_\epsilon(x^{(i)}) = \frac{1 + \epsilon_i \gamma}{2}.$$

For the cost class \mathcal{G} corresponds to \mathcal{V}' , we consider its subset \mathcal{U}_γ :

$$\mathcal{U}_\gamma = \{g_\epsilon(x, a) : g_\epsilon(x, 0) = h_\epsilon(x), \quad g_\epsilon(x, 1) = \frac{1}{2}, \epsilon \in \{-1, 1\}^m\}.$$

Let P^γ be the uniform distribution over $\{x^{(1)}, \dots, x^{(m)}\}$, then for

$$\pi_\epsilon^*(x^{(i)}) = \begin{cases} 0, & \text{if } \epsilon_i = -1, \\ 1, & \text{otherwise.} \end{cases},$$

we have

$$\max_a g_\epsilon(x, a) - \min_a g_\epsilon(x, a) > \frac{\gamma}{2}, \quad \frac{1}{4} \leq \frac{1-\gamma}{2} \leq \mathbb{E}_x[\min_a g_\epsilon(x, a)] < \mathbb{E}_x[\max_a g_\epsilon(x, a)] \leq \frac{1+\gamma}{2} \leq \frac{3}{4}.$$

Now for every sufficiently large T , the argument in Foster and Rakhlin (2020) shows that when letting $\gamma \propto T^{-\frac{1}{2+p_2}}/\text{polylog}(T)$, $P_{\mathcal{X}} = P_\gamma$, for every every policy π , there exists some ϵ_0 so that when the reward is generated as $r(x_t, a) \sim \text{Bernoulli}(g_{\epsilon_0}(x, a))$, it holds that

$$\frac{\gamma}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbf{1} \{ \pi_{\epsilon_0}^*(x_t) \neq \pi_t(x_t) \} \right] = \mathbb{E} \left[\sum_{t=1}^T g_{\epsilon_0}(x_t, \pi_{\epsilon_0}^*(x_t)) - g_{\epsilon_0}(x_t, \pi_t(x_t)) \right] \gtrsim \gamma T.$$

This implies the constructed \mathcal{G} is $\tilde{\Theta}(T^{-\frac{1}{2+p_2}})$ -inseparable.

Step2: Using the same proof as Theorem 2 in Foster and Rakhlin (2020), for \mathcal{W} with $H(\mathcal{W}, \varepsilon) = \Theta(\varepsilon^{-p_1})$, we can construct some \mathcal{W}' so that the corresponding \mathcal{F}' satisfies $\text{Reg}_{\text{CB}}(T; \mathcal{F}) = \tilde{\Omega}(T^{\frac{1+p_1}{2+p_1}})$, we just construct \mathcal{F} by adding a constant function $f(x, a) \equiv 1$ into \mathcal{F}' .

Step3: By our construction, we have there exists some $f_0 \in \mathcal{F}$ so that $f(x, a) \equiv 1$, $\text{Reg}_{\text{CB}}(T; \mathcal{F}) = \tilde{\Omega}(T^{\frac{1+p_1}{2+p_1}})$. And \mathcal{G} is $\tilde{\Theta}(T^{-\frac{1}{2+p_2}})$ -inseparable, applying Theorem 3.2 leads to the $\tilde{\Omega}(\max\{T^{\frac{1+p_1}{2+p_1}}, T^{\frac{1+p_1}{2+p_2}}\}) = \tilde{\Omega}(T^{\frac{1+p_1}{2+p_1}} + T^{\frac{1+p_1}{2+p_2}})$ lower bound as desired. \square

B PROOF OF RESULTS IN SECTION 4

B.1 Estimation Error of Online-To-Batch Conversion Oracle

Given online regression oracles $\mathcal{R}^r, \mathcal{R}^c$ satisfying (2) and (3), we formally define the OTB oracles $\mathcal{R}_{\text{est}}^r, \mathcal{R}_{\text{est}}^c$ as following: Given the dataset \mathcal{D}_a as in Assumption 4.1, if we run online oracles $\mathcal{R}^r, \mathcal{R}^c$ over \mathcal{D}_a , and suppose the output of online oracles are given by $\{f_i(\cdot, a)\}_{i=1}^M, \{g_i(\cdot, a)\}_{i=1}^M$, the output $\hat{f}(\cdot, a), \hat{g}(\cdot, a)$ of $\mathcal{R}_{\text{est}}^r, \mathcal{R}_{\text{est}}^c$ is defined as following:

$$\hat{f}(x, a) := \frac{1}{M} \sum_{i=1}^M f_i(x, a), \quad \hat{g}(x, a) := \frac{1}{M} \sum_{i=1}^M g_i(x, a). \quad (8)$$

Now we would show the following estimation error guarantee of \hat{f}, \hat{g} defined above:

Lemma B.1. *For every a , suppose the online regression oracles $\mathcal{R}^r, \mathcal{R}^c$ satisfies (2) and (3), we have $\hat{f}(\cdot, a), \hat{g}(\cdot, a)$ defined in (8) satisfies Assumption 4.1 with*

$$\mathcal{E}_\delta(M; \mathcal{F}) \lesssim \frac{\text{Reg}_{\text{sq}}^r(M) \log(1/\delta)}{M}, \quad \mathcal{E}_\delta(M; \mathcal{G}) \lesssim \frac{\text{Reg}_{\text{sq}}^c(M) \log(1/\delta)}{M}.$$

Proof of Lemma B.1. We would prove the result for \hat{g} , the proof of \hat{f} is similar.

By the convexity of $x \rightarrow x^2$, we have

$$\begin{aligned} & \mathbb{E}_x \left[\left\| \frac{1}{M} \sum_{t=1}^M \mathbf{g}_t(x, a) - \mathbf{g}^*(x, a) \right\|_\infty^2 \right] \\ & \leq \frac{1}{M} \sum_{t=1}^M \mathbb{E}_x \left[\left\| \mathbf{g}_t(x, a) - \mathbf{g}^*(x, a) \right\|_\infty^2 \right] \\ & \leq \frac{\text{Reg}_{\text{sq}}^r(M)}{M} + \frac{1}{M} \sum_{t=1}^M \left(\mathbb{E}_x \left[\left\| \mathbf{g}_t(x, a) - \mathbf{g}^*(x, a) \right\|_\infty^2 \right] - \left\| \mathbf{g}_t(x_t, a) - \mathbf{g}^*(x_t, a) \right\|_\infty^2 \right). \end{aligned}$$

In particular, we have for $W_t := \underbrace{\left\| \mathbf{g}_t(x_t, a) - \mathbf{g}^*(x_t, a) \right\|_\infty^2}_{:=V_t} - \mathbb{E}_x \left[\left\| \mathbf{g}_t(x, a) - \mathbf{g}^*(x, a) \right\|_\infty^2 \right]$, $\mathbb{E}[W_t | \mathcal{F}'_{t-1}] = 0$ and $|V_t| \leq 1$.

As a result,

$$\text{Var}(W_t | \mathcal{F}'_{t-1}) \leq \mathbb{E}[W_t^2 | \mathcal{F}'_{t-1}] = \text{Var}[V_t | \mathcal{F}'_{t-1}] \leq \mathbb{E}[V_t^2 | \mathcal{F}'_{t-1}]$$

$$\leq \mathbb{E}[V_t | \mathcal{F}'_{t-1}] = \mathbb{E}_x[\|\mathbf{g}_t(x, a) - \mathbf{g}^*(x, a)\|_\infty^2]$$

Now by Lemma A.1, we have with probability at least $1 - 2\delta \log T$,

$$\begin{aligned} & \left| \sum_{t=1}^M \|\mathbf{g}_t(x_t, a) - \mathbf{g}^*(x_t, a)\|_\infty^2 - \mathbb{E}_x[\|\mathbf{g}_t(x, a) - \mathbf{g}^*(x, a)\|_\infty^2] \right| = \left| \sum_{t=1}^M W_t \right| \\ & \lesssim \sqrt{\log(1/\delta)} \sqrt{\max\left\{ \sum_{t=1}^M \mathbb{E}_x[\|\mathbf{g}_t(x, a) - \mathbf{g}^*(x, a)\|_\infty^2], \log(1/\delta) \right\}} \end{aligned}$$

Thus when $\max\{\sum_{t=1}^M \mathbb{E}_x[\|\mathbf{g}_t(x, a) - \mathbf{g}^*(x, a)\|_\infty^2] > \log(1/\delta)$, we have

$$\sum_{t=1}^M \mathbb{E}_x[\|\mathbf{g}_t(x, a) - \mathbf{g}^*(x, a)\|_\infty^2] \lesssim \text{Reg}_{\text{sq}}^c(M) + \sqrt{\log(1/\delta) \sum_{t=1}^M \mathbb{E}_x[\|\mathbf{g}_t(x, a) - \mathbf{g}^*(x, a)\|_\infty^2]},$$

which then implies

$$\sum_{t=1}^M \mathbb{E}_x[\|\mathbf{g}_t(x, a) - \mathbf{g}^*(x, a)\|_\infty^2] \lesssim \text{Reg}_{\text{sq}}^c(M) + \log(1/\delta).$$

That leads to the desired result with replacing δ by $\frac{\delta}{2 \log T}$. □

B.2 Proof of Theorem 4.1

Proof of Theorem 4.1. Since in Algorithm 2, we use $O(KT_0)$ steps to take exploration, the expected regret incurred by this stage is upper bounded by $O(\frac{T_{\text{OPT}}}{B} + 1)KT_0$. For the second stage, since $Z \leq CT_{\text{OPT}}/B$, the expected regret is upper bounded by

$$\left(\frac{T_{\text{OPT}}}{B} + 1\right) \sqrt{KT(\text{Reg}_{\text{sq}}^r(T) + \text{Reg}_{\text{sq}}^c(T) + \log T)}$$

Then Theorem 4.1 holds by add the regret in these two stages together. □

B.3 Proof of Lemma 4.1

Proof of Lemma 4.1. Similar to Agrawal and Devanur (2016) in linear case, our proof relies on the result about the ‘‘intermediate sample optimal’’ $\overline{\text{OPT}}^\epsilon$, which is the value of

$$\begin{aligned} & \max_{\mathbf{p} \in (\Delta^K)^{T_0}} \frac{1}{T_0} \sum_{t=KT_0+1}^{(K+1)T_0} f^*(x_t, a) p_a(x_t) \\ & \text{subject to } \frac{1}{T_0} \sum_{t=KT_0+1}^{(K+1)T_0} \mathbf{g}^*(x_t, a) p_a(x_t) \leq \mathbf{B}/T + \epsilon \mathbf{1}. \end{aligned} \tag{9}$$

Applying the Lemma F.4 and F.6 of Agrawal and Devanur (2014b) in the same way as Agrawal and Devanur (2016)¹ leads to

$$\text{OPT} - \epsilon \leq \overline{\text{OPT}}^\epsilon \leq \text{OPT} + 2\epsilon \left(\frac{T_{\text{OPT}}}{B} + 1\right) \tag{10}$$

with probability $1 - O(1/T^2)$ when $\epsilon \geq 2\sqrt{\frac{\log(T_0 d)}{T_0}}$.

¹Notice that our definition of ϵ and OPT corresponds to the γ/T and OPT/T in Agrawal and Devanur (2016).

Now it sufficient to bound $\widehat{\text{OPT}}(T_0)$ and $\overline{\text{OPT}}^\epsilon$: Firstly noticing that since we have used T_0 i.i.d. samples from P_X to estimate $f(\cdot, a), \mathbf{g}(\cdot, a)$, we have with probability at least $1 - O(1/T^2)$,

$$\mathbb{E}_x[(\hat{f}_0(x, a) - f(x, a))^2] \leq \mathcal{E}_{\delta_0}(T_0; \mathcal{F}), \quad \mathbb{E}_x[\|\hat{\mathbf{g}}_0(x, a) - \mathbf{g}(x, a)\|_\infty^2] \leq d\mathcal{E}_{\delta_0}(T_0; \mathcal{G}), \quad \forall a \in [K], \quad (11)$$

where $\delta_0 = 1/TK^2$. Thus we have with probability at least $1 - O(1/T^2)$, for all $\mathbf{p} \in (\Delta^K)^{T_0}$,

$$\begin{aligned} \left| \sum_{t=KT_0+1}^{(K+1)T_0} \sum_a p_a(x_t) [\hat{f}_0(x, a) - f^*(x, a)] \right| &\leq \sum_{t=KT_0+1}^{(K+1)T_0} \max_a |\hat{f}_0(x, a) - f^*(x, a)| \\ &\leq \sum_{t=KT_0+1}^{(K+1)T_0} \sqrt{\mathbb{E}[\max_a |\hat{f}_0(x, a) - f^*(x, a)|^2]} + 4\sqrt{T_0 \log T}, \\ &\leq T_0 \sqrt{K\mathcal{E}_{\delta_0}(T_0; \mathcal{F})} + 4\sqrt{T_0 \log T} \end{aligned}$$

where the second line is by Cauchy-Schwartz inequality and Hoeffding's inequality and the third line is by (11). Applying the same argument to $\hat{\mathbf{g}}_0$, we get

$$\mathbb{P}\left(\left| \sum_{t=KT_0+1}^{(K+1)T_0} \sum_a p_a(x_t) [\hat{f}_0(x, a) - f^*(x, a)] \right| \geq \underbrace{T_0 \sqrt{K\mathcal{E}_{\delta_0}(T_0; \mathcal{F})} + 4\sqrt{T_0 \log T}}_{\leq T_0 \mathcal{R}(T_0)}, \forall \mathbf{p} \in (\Delta^K)^{T_0}\right) \lesssim 1/T^2.$$

and

$$\mathbb{P}\left(\left\| \sum_{t=KT_0+1}^{(K+1)T_0} \sum_a p_a(x_t) [\hat{\mathbf{g}}_0(x, a) - \mathbf{g}^*(x, a)] \right\|_\infty \geq \underbrace{(T_0 \sqrt{K\mathcal{E}_{\delta_0}(T_0; \mathcal{G})} + 4\sqrt{T_0 \log T})}_{\leq T_0 \mathcal{R}(T_0)}, \forall \mathbf{p} \in (\Delta^K)^{T_0}\right) \lesssim 1/T^2.$$

Denoting $\bar{\mathbf{p}}$ the optimal solution of (9) for $2\sqrt{\frac{\log(Td)}{T_0}} \leq \epsilon \leq \mathcal{R}(T_0)$, we have the deviation bound of $\hat{\mathbf{g}}_0$ implies $\bar{\mathbf{p}}$ is feasible in (6) with probability $1 - O(1/T^2)$, thus combining the deviation bound of \hat{f}_0 , we have with probability at least $1 - O(1/T^2)$,

$$\widehat{\text{OPT}}(T_0) \geq \frac{1}{T_0} \sum_{t=KT_0+1}^{(K+1)T_0} \sum_a \bar{p}_a(x_t) \hat{f}_0(x_t, a) \geq \frac{1}{T_0} \sum_{t=KT_0+1}^{(K+1)T_0} \sum_a \bar{p}_a(x_t) f^*(x_t, a) - \mathcal{R}(T_0) = \overline{\text{OPT}}^\epsilon - \mathcal{R}(T_0). \quad (12)$$

On the other hand, for $\hat{\mathbf{p}}$ the optimal solution of (6), we have with high probability $\hat{\mathbf{p}}$ is feasible for (9) with $\epsilon_0 = 2\mathcal{R}(T_0)$, thus with probability at least $1 - O(1/T^2)$, we have

$$\widehat{\text{OPT}}(T_0) = \frac{1}{T_0} \sum_{t=KT_0+1}^{(K+1)T_0} \sum_a \hat{p}_a(x_t) \hat{f}_0(x_t, a) \leq \frac{1}{T_0} \sum_{t=KT_0+1}^{(K+1)T_0} \sum_a \hat{p}_a(x_t) f_0^*(x_t, a) + \mathcal{R}(T_0) \leq \overline{\text{OPT}}^{\epsilon_0} + \mathcal{R}(T_0). \quad (13)$$

Now combine (10),(12),(13), we get with probability at least $1 - O(1/T^2)$,

$$\text{OPT} \leq \widehat{\text{OPT}}(T_0) + \mathcal{R}(T_0) \leq \text{OPT} + 6\mathcal{R}(T_0) \left(\frac{T\text{OPT}}{B} + 1\right),$$

as desired. \square

B.4 Result in Linear CBwK

In linear CBwK setting, suppose both \mathcal{F} and $\tilde{\mathcal{G}}$ are m -dimensional linear classes. If we select the online oracles $\mathcal{R}^r, \tilde{\mathcal{R}}^c$ as in Table 1 and offline oracles $\mathcal{R}_{est}^r, \mathcal{R}_{est}^c$ as the OTB oracle constructed from $\mathcal{R}^r, \mathcal{R}^c$, we have then

$$\text{Reg}_{\text{sq}}^r(T) \lesssim m \log T, \text{Reg}_{\text{sq}}^c(T) \lesssim md \log T, \mathcal{E}_{T_0}(\mathcal{F}) \lesssim \frac{m \log T}{T_0}, \mathcal{E}_{T_0}(\mathcal{G}) \lesssim \frac{md \log T}{T_0}. \quad (14)$$

Now bringing (14) to Theorem 4.1 and selecting $T_0 = (md)^{1/3} \sqrt{T/K}$ leads to

$$\text{Reg}(T) \lesssim \left(\frac{T\text{OPT}}{B} + 1\right) (\sqrt{KTdm})$$

when $B = \tilde{\Omega}((md)^{1/3} \sqrt{KT})$.

B.5 Results in Nonparametric CBwK

In nonparametric CBwK setting, suppose both \mathcal{F} and $\tilde{\mathcal{G}}$ are p -nonparametric classes. If we select the online oracles $\mathcal{R}^r, \tilde{\mathcal{R}}^c$ as in Table 1 and offline oracles $\mathcal{R}_{est}^r, \mathcal{R}_{est}^c$ as the OTB oracle constructed from $\mathcal{R}^r, \mathcal{R}^c$, we have then

$$\text{Reg}_{\text{sq}}^r(T) \lesssim (KT)^{1-\frac{2}{2+p}}, \text{Reg}_{\text{sq}}^c(T) \lesssim d(KT)^{1-\frac{2}{2+p}}, \mathcal{E}_{T_0}(\mathcal{F}) \lesssim \frac{(KT_0)^{1-\frac{2}{2+p}}}{T_0}, \mathcal{E}_{T_0}(\mathcal{G}) \lesssim \frac{d(KT)^{1-\frac{2}{2+p}}}{T_0}. \quad (15)$$

Now bringing (15) to Theorem 4.1 and selecting $T_0 = d^{\frac{2+p}{6+2p}} K^{\frac{-1}{2+p}} T^{\frac{1+p}{2+p}}$ leads to

$$\text{Reg}(T) \lesssim \tilde{O}\left(\left(\frac{TOPT}{B} + 1\right)\sqrt{d}(KT)^{\frac{1+p}{2+p}}\right)$$

when $B = \tilde{\Omega}(d^{\frac{2+p}{6+2p}} (KT)^{\frac{3+p}{4+2p}})$.

C NUMERICAL RESULTS

We provide simulation results in this section to validate the time horizon (T), dimension (m), and number of arms (K) dependencies of SquareCBwK, which employs Newtonized GLMtron and Online Gradient Descent oracles, for the linear CBwK. We also conduct a performance comparison of SquareCBwK with LinUCB (Agrawal and Devanur, 2016).

For general distributions of $x_{t,a}$, solving the OPT value from population linear programming (1) is challenging which renders the computation of regret, $TOPT - \mathbb{E}[\sum_{t=1}^T r_{t,a_t}]$, an intractable task. Therefore, in our subsequent simulations, we adhere to the fixed context setting, which simplifies the process of determining the OPT value to that of solving a linear programming, as outlined in Section C.1. In fact, the fixed context setting is a special case of i.i.d. context by letting the distribution of contexts be the point mass distribution.

C.1 Experiment Setting

Throughout the experiment, we assume the linear structured reward and cost classes. For any fixed dimension m , arm number $K \leq m - 1$, and number of constraints $4 \leq d \leq m - 1$, we set

1. Underlying parameters:

$$\theta_0 = \frac{1}{\sqrt{2}}(e_1 + e_2), \theta_1 = \frac{1}{\sqrt{2}}(e_1 + e_3), \theta_2 = \frac{1}{2}(e_2 + e_3 + e_4 + e_5), \theta_i = e_{i+1}, 3 \leq i \leq d.$$

2. **Fixed context set:** At every round t , the contexts $x_{t,a}$ are given by $x_{t,a} = \frac{1}{\sqrt{2}}e_1 + e_{a+1}, a \in [K]$

3. **Generation of rewards and costs:** At every round t , after an action $a_t \in [K]$ is selected

$$r_t = \langle x_{t,a_t}, \theta_0 \rangle + \epsilon_{t,0}, \quad c_{t,i} = \langle x_{t,a_t}, \theta_i \rangle + \epsilon_{t,i}, i \in [d]$$

with $\epsilon_{t,i} \sim_{i.i.d.} \mathcal{N}(0, 0.2)$.

In the experiments that follow, we will simulate the three algorithms in Table 3 independently with one varying hyper-parameter, selected from m, K , and T , while keeping the remaining two hyper-parameters constant. This will allow us to validate the theoretical dependency as detailed in Table 3. The source code for reproducing these results can be found in <https://github.com/quejialin/SquareCBwK>

Parameter	GLMtronNewton	Online GD	LinUCB
K	$O(\sqrt{K})$	$O(\sqrt{K})$	$O(1)$
m	$O(\sqrt{m})$	$O(1)$	$O(m)$
T	$O(\sqrt{T})$	$O(T^{3/4})$	$O(\sqrt{T})$

Table 3: Theoretical dependency of SquareCBwK with Newtonized GLMtron oracle, SquareCBwK with Online Gradient Descent oracles, and LinUCB algorithms (Agrawal and Devanur, 2016) on K, m, T

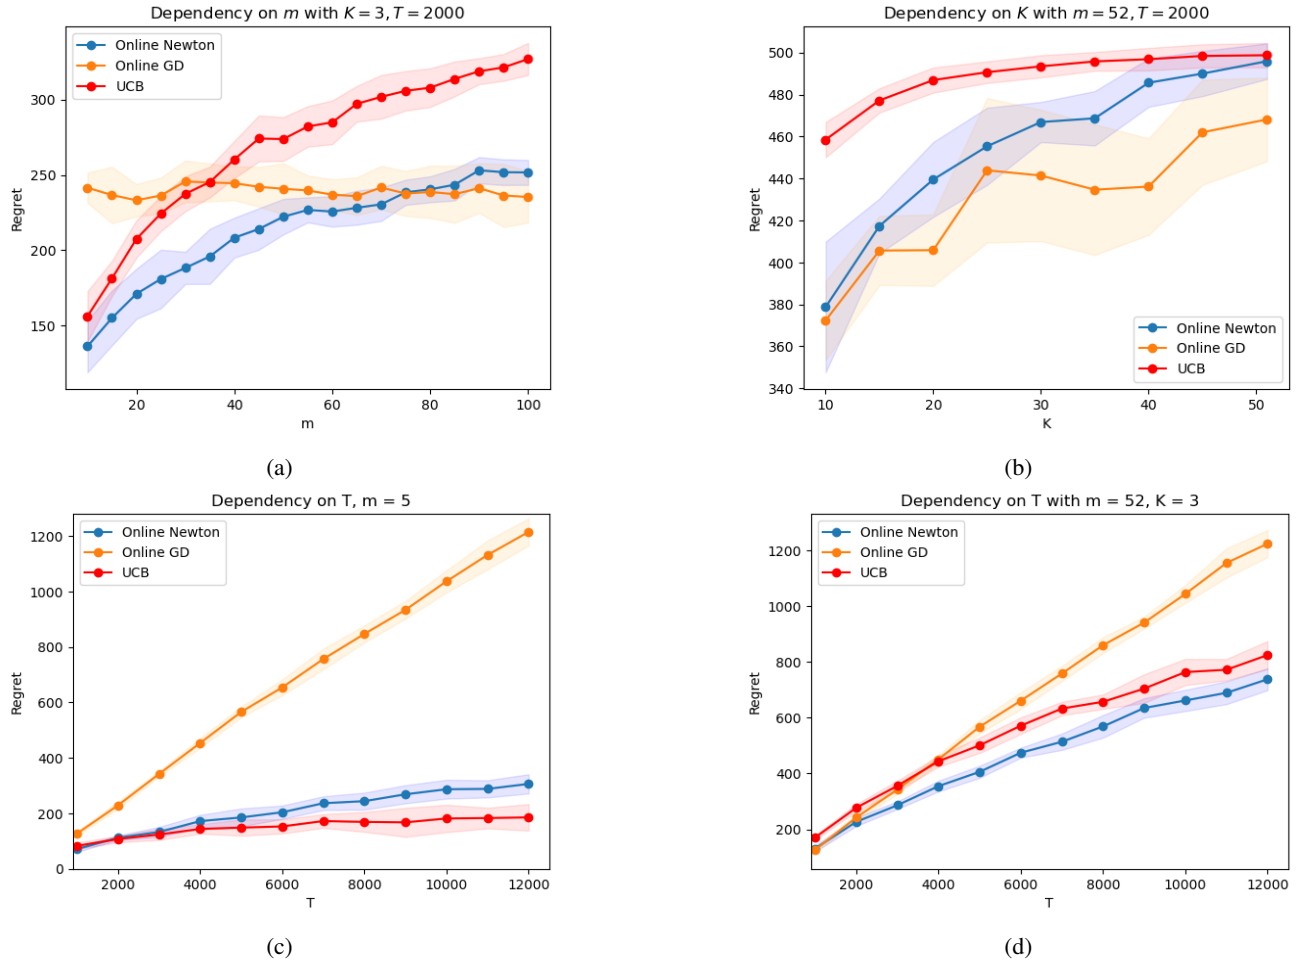


Figure 1: Dependency of different algorithms on different parameters m, K, T , where the blue curve corresponds to the regret of SquareCBwK with Newtonized GLMtron oracle, the orange curve shows regret performance of SquareCBwK with online Gradient oracle, and the red curve represents the regret of LinUCB.

Dependency on m

To examine the dependency on m , we fix $K = 3, T = 2000$ and simulate different algorithms with $m \in \{10, 14, 19, \dots, 96, 101\}$ over 10 times. The regret curve is presented in Figure 1(a), from which we can see that the regret of the SquareCBwK with Online-GD oracle is almost unaffected by m , and that the regret of GLMtron oracle grows at a slower pace than LinUCB as m increases, matching the theoretical guarantees in Table 3.

Dependency on K

To examine the dependency on K , we fix $m = 52, T = 2000$ and simulate different algorithms with $K \in \{5, 10, \dots, 45, 50\}$ over 10 times. Note that here we pick a large $m = 52$ since we need to make sure the condition $K \leq m - 1$ holds. The regret curve is presented in Figure 1(b). Although the dependency of LinUCB on K is better than SquareCBwK with GLMtron or Online-GD oracles, the regret of LinUCB is larger in the large- m regime.

Dependency on T

To examine the dependency on T , we fix $K = 3$ and simulate different algorithms with $T \in \{1000, 1100, \dots, 12000\}$. Here we also choose $m \in \{5, 52\}$ to compare the influence of large and small dimension respectively. The regret curves are presented in Figure 1(c) and Figure 1(d). It can be shown that the regret of SquareCBwK with Online-GD grows much faster than SquareCBwK with GLMtron and LinUCB for both m as T increases. The regret of GLMtron and LinUCB are of

the same order. Although LinUCB outperforms SquareCBwK with GLMtron slightly in the small m regime, SquareCBwK with GLMtron exhibits an improved performance in the large m regime, which again verifies the dependency on T and m in Table 3 of different algorithms.

From the above comparisons, we can see that in the large dimension setting of linear CBwK, SquareCBwK with online-GD oracle is independent of m although with worse dependency on T , and SquareCBwK with Newtonized GLMtron oracles exhibits a better dependency on dimension m than LinUCB, while keeping the same order of T as LinUCB. These results together verify the superiority of our algorithm in the large dimension setting of linear CBwK.