# Linear Contextual Bandits with Knapsacks 

Shipra Agrawal*<br>Nikhil R. Devanur ${ }^{\dagger}$


#### Abstract

We consider the linear contextual bandit problem with resource consumption, in addition to reward generation. In each round, the outcome of pulling an arm is a reward as well as a vector of resource consumptions. The expected values of these outcomes depend linearly on the context of that arm. The budget/capacity constraints require that the total consumption doesn't exceed the budget for each resource. The objective is once again to maximize the total reward. This problem turns out to be a common generalization of classic linear contextual bandits (linContextual) $[8,11,1]$, bandits with knapsacks (BwK) [3, 9], and the online stochastic packing problem (OSPP) [4, 14]. We present algorithms with near-optimal regret bounds for this problem. Our bounds compare favorably to results on the unstructured version of the problem $[5,10]$ where the relation between the contexts and the outcomes could be arbitrary, but the algorithm only competes against a fixed set of policies accessible through an optimization oracle. We combine techniques from the work on linContextual, BwK and OSPP in a nontrivial manner while also tackling new difficulties that are not present in any of these special cases.


## 1 Introduction

In the contextual bandit problem [8,2], the decision maker observes a sequence of contexts (or features). In every round she needs to pull one out of $K$ arms, after observing the context for that round. The outcome of pulling an arm may be used along with the contexts to decide future arms. Contextual bandit problems have found many useful applications such as online recommendation systems, online advertising, and clinical trials, where the decision in every round needs to be customized to the features of the user being served. The linear contextual bandit problem $[1,8,11]$ is a special case of the contextual bandit problem, where the outcome is linear in the feature vector encoding the context. As pointed by [2], contextual bandit problems represent a natural half-way point between supervised learning and reinforcement learning: the use of features to encode contexts and the models for the relation between these feature vectors and the outcome are often inherited from supervised learning, while managing the exploration-exploitation tradeoff is necessary to ensure good performance in reinforcement learning. The linear contextual bandit problem can thus be thought of as a midway between the linear regression model of supervised learning, and reinforcement learning.

Recently, there has been a significant interest in introducing multiple "global constraints" in the standard bandit setting [9, 3, 10, 5]. Such constraints are crucial for many important real-world applications. For example, in clinical trials, the treatment plans may be constrained by the total availability of medical facilities, drugs and other resources. In online advertising, there are budget constraints that restrict the number of times an ad is shown. Other applications include dynamic pricing, dynamic procurement, crowdsourcing, etc.; see [9, 3] for many such examples.

In this paper, we consider the linear contextual bandit with knapsacks (henceforth, linCBwK) problem. In this problem, the context vectors are generated i.i.d. in every round from some unknown distribution, and on picking an arm, a reward and a consumption vector is observed, which depend

[^0]linearly on the context vector. The aim of the decision maker is to maximize the total reward while ensuring that the total consumption of every resource remains within a given budget. Below, we give a more precise definition of this problem. We use the following notational convention throughout: vectors are denoted by bold face lower case letters, while matrices are denoted by regular face upper case letters. Other quantities such as sets, scalars, etc. may be of either case, but never bold faced. All vectors are column vectors, i.e., a vector in $n$ dimensions is treated as an $n \times 1$ matrix. The transpose of matrix $A$ is $A^{\top}$.

Definition 1 (linCBwK). There are $K$ "arms", which we identify with the set $[K]$. The algorithm is initially given as input a budget $B \in \mathbb{R}_{+}$. In every round $t$, the algorithm first observes context $\mathbf{x}_{t}(a) \in[0,1]^{m}$ for every arm $a$, and then chooses an arm $a_{t} \in[K]$, and finally observes a reward $r_{t}\left(a_{t}\right) \in[0,1]$ and a d-dimensional consumption vector $\mathbf{v}_{t}\left(a_{t}\right) \in[0,1]^{d}$. The algorithm has a "no-op" option, which is to pick none of the arms and get 0 reward and $\mathbf{0}$ consumption. The goal of the algorithm is to pick arms such that the total reward $\sum_{t=1}^{T} r_{t}\left(a_{t}\right)$ is maximized, while ensuring that the total consumption does not exceed the budget, i.e., $\sum_{t} \mathbf{v}_{t}\left(a_{t}\right) \leq B 1$.
We make the following stochastic assumption for context, reward, and consumption vectors. In every round $t$, the tuple $\left\{x_{t}(a), r_{t}(a), \mathbf{v}_{t}(a)\right\}_{a=1}^{K}$ is generated from an unknown distribution $\mathcal{D}$, independent of everything in previous rounds. Also, there exists an unknown vector $\mu_{*} \in[0,1]^{m}$ and a matrix $W_{*} \in[0,1]^{m \times d}$ such that for every arm a, given contexts $x_{t}(a)$, and history $H_{t-1}$ before time $t$,

$$
\begin{equation*}
\mathbb{E}\left[r_{t}(a) \mid x_{t}(a), H_{t-1}\right]=\mu_{*}^{\top} x_{t}(a), \quad \mathbb{E}\left[\mathbf{v}_{t}(a) \mid x_{t}(a), H_{t-1}\right]=W_{*}^{\top} x_{t}(a) . \tag{1}
\end{equation*}
$$

For succinctness, we will denote the tuple of contexts for $K$ arms at time $t$ as matrix $X_{t} \in[0,1]^{m \times K}$, with $\mathbf{x}_{t}(a)$ being the $a^{\text {th }}$ column of this matrix. Similarly, rewards and consumption vectors at time $t$ are represented as the vector $\mathbf{r}_{t} \in[0,1]^{K}$ and the matrix $V_{t} \in[0,1]^{d \times K}$ respectively.

As we discuss later in the text, the assumption in equation (1) forms the primary distinction between our linear contextual bandit setting and the general contextual bandit setting considered in [5]. Exploiting this linearity assumption will allow us to generate regret bounds which do not depend on the number of arms $K$, rendering it to be especially useful when the number of arms is large. Some examples of this include recommendation systems with large number of products (e.g., retail products, travel packages, ad creatives, sponsored facebook posts). Another advantage over using the general contextual bandit setting of [5] is that we don't need an oracle access to a certain optimization problem, which in this case is required to solve an NP-Hard problem. (See Section 1.1 for a more detailed discussion.)
We compare the performance of an algorithm to that of an optimal adaptive policy that knows the distribution $\mathcal{D}$ and the parameters $\left(\mu_{*}, W_{*}\right)$, and can take into account the history up to that point, as well as the current context, to decide (possibly with randomization) which arm to pull at time $t$. However, it is easier to work with an upper bound on this, which is the optimal expected reward of a static policy that is required to satisfy the constraints only in expectation. This technique has been used in several related problems and is standard by now [14, 9].
Definition 2 (Optimal Static Policy). A context-dependent non-adaptive policy $\pi$ is a mapping from context space $[0,1]^{m \times K}$ to $\Omega=\left\{\mathbf{p} \in[0,1]^{K}:\|\mathbf{p}\|_{1} \leq 1\right\}$, where $\pi(X)_{i}$ denotes the probability of playing arm $i$ when the context is $X$, and $1-\sum_{i=1}^{K} \pi(X)_{i}$ is the probability of no-op. Define $\mathbf{r}(\pi)$ and $\mathbf{v}(\pi)$ to be the expected reward and consumption vector of policy $\pi$, respectively, i.e.

$$
\begin{align*}
\mathbf{r}(\pi) & :=\mathbb{E}_{(X, \mathbf{r}, V) \sim \mathcal{D}}[\mathbf{r} \pi(X)]=\mathbb{E}_{X \sim \mathcal{D}}\left[\mu_{*}^{\top} X \pi(X)\right] .  \tag{2}\\
\mathbf{v}(\pi) & :=\mathbb{E}_{(X, \mathbf{r}, V) \sim \mathcal{D}}[V \pi(X)]=\mathbb{E}_{X \sim \mathcal{D}}\left[W_{*}^{\top} X \pi(X)\right] .  \tag{3}\\
\text { Let } \pi^{*} & :=\arg \max _{\pi} \quad T \mathbf{r}(\pi) \text { such that } T \mathbf{v}(\pi) \leq B \mathbf{1} \tag{4}
\end{align*}
$$

be the optimal static policy. Note that since no-op is allowed, a feasible policy always exists. We denote the value of this optimal static policy by OPT $:=T \mathbf{r}\left(\pi^{*}\right)$.

The following lemma proves that OPT upper bounds the value of an optimal adaptive policy. Proof is in Appendix B in the supplement.
Lemma 1. Let $\overline{O P T}$ denote the value of an optimal adaptive policy that knows the distribution $\mathcal{D}$ and parameters $\mu_{*}, W_{*}$. Then $O P T \geq \overline{O P T}$.

Definition 3 (Regret). Let $a_{t}$ be the arm played at time $t$ by the algorithm. Then, regret is defined as

$$
\operatorname{regret}(T):=O P T-\sum_{t=1}^{T} \mathbf{r}_{t}\left(a_{t}\right)
$$

### 1.1 Main results

Our main result is an algorithm with near-optimal regret bound for linCBwK.
Theorem 1. There is an algorithm for linCBwK such that if $B>m^{1 / 2} T^{3 / 4}$, then with probability at least $1-\delta$,

$$
\operatorname{regret}(T)=O\left(\left(\frac{O P T}{B}+1\right) m \sqrt{T \log (d T / \delta) \log (T)}\right)
$$

Relation to general contextual bandits. There have been recent papers [5, 10] that solve problems similar to linCBwK but for general contextual bandits. In these papers the relation between contexts and outcome vectors is arbitrary and the algorithms compete with an arbitrary fixed set of context dependent policies $\Pi$ accessible via an optimization oracle, with regret bounds being $O\left(\left(\frac{\mathrm{OPT}}{B}+1\right) \sqrt{K T \log (d T|\Pi| / \delta)}\right)$. These approaches could potentially be applied to the linear setting using a set $\Pi$ of linear context dependent policies. Comparing their bounds with ours, in our results, essentially a $\sqrt{K \log (|\Pi|)}$ factor is replaced by a factor of $m$. Most importantly, we have no dependence on $K,{ }^{3}$ which enables us to consider problems with large action spaces.
Further, suppose that we want to use their result with the set of linear policies, i.e., policies of the form, for some fixed $\boldsymbol{\theta} \in \Re^{m}$,

$$
\arg \max _{a \in[K]}\left\{\mathbf{x}_{t}(a)^{\top} \boldsymbol{\theta}\right\} .
$$

Then, their algorithms would require access to an "Arg-Max Oracle" that can find the best such policy (maximizing total reward) for a given set of contexts and rewards (no resource consumption). In fact, by a reduction from the problem of learning halfspaces with noise [16], we can show that the optimization problem underlying such an "Arg-Max Oracle" problem is NP-Hard, making such an approach computationally expensive. The proof of this is in Appendix C in the supplement.

The only downside to our results is that we need the budget $B$ to be $\Omega\left(m^{1 / 2} T^{3 / 4}\right)$. Getting similar bounds for budgets as small as $B=\Theta(m \sqrt{T})$ is an interesting open problem. (This also indicates that this is indeed a harder problem than all the special cases.)

Near-optimality of regret bounds. In [12], it was shown that for the linear contextual bandits problem, no online algorithm can achieve a regret bound better than $\Omega(m \sqrt{T})$. In fact, they prove this lower bound for linear contextual bandits with static contexts. Since that problem is a special case of the linCBwK problem with $d=1$, this shows that the dependence on $m$ and $T$ in the above regret bound is optimal upto $\log$ factors. For general contextual bandits with resource constraints, the bounds of $[5,10]$ are near optimal.

Relation to BwK [3] and OSPP [4]. It is easy to see that the linCBwK problem is a generalization of the linear contextual bandits problem [1, 8, 11]. There, the outcome is scalar and the goal is to simply maximize the sum of these. Remarkably, the linCBwK problem also turns out to be a common generalization of the bandits with knapsacks (BwK) problem considered in [9, 3], and the online stochastic packing problem (OSPP) studied by [13, 6, 15, 14, 4]. In both BwK and OSPP, the outcome of every round $t$ is a reward $r_{t}$ and a vector $\mathbf{v}_{t}$ and the goal of the algorithm is to maximize $\sum_{t=1}^{T} r_{t}$ while ensuring that $\sum_{t=1}^{T} \mathbf{v}_{t} \leq B 1$. The problems differ in how these rewards and vectors are picked. In the OSPP problem, in every round $t$, the algorithm may pick any reward, vector pair from a given set $A_{t}$ of $d+1$-dimensional vectors. The set $A_{t}$ is drawn i.i.d. from an unknown distribution over sets of vectors. This corresponds to the special case of linCBwK, where $m=d+1$ and the context $\mathbf{x}_{t}(a)$ itself is equal to $\left(r_{t}(a), \mathbf{v}_{t}(a)\right)$. In the BwK problem, there is a fixed set of arms, and for each arm there is an unknown distribution over reward,vector pairs. The algorithm picks an arm and a reward, vector pair is drawn from the corresponding distribution for that arm. This

[^1]corresponds to the special case of linCBwK, where $m=K$ and the context $X_{t}=I$, the identity matrix, for all $t$.

We use techniques from all three special cases: our algorithms follow the primal-dual paradigm and use an online learning algorithm to search the dual space, as was done in [3]. In order to deal with linear contexts, we use techniques from $[1,8,11]$ to estimate the weight matrix $W_{*}$, and define "optimistic estimates" of $W_{*}$. We also use the technique of combining the objective and the constraints using a certain tradeoff parameter and that was introduced in [4]. Further new difficulties arise, such as in estimating the optimum value from the first few rounds, a task that follows from standard techniques in each of the special cases but is very challenging here. We develop a new way of exploration that uses the linear structure, so that one can evaluate all possible choices that could have led to an optimum solution on the historic sample. This technique might be of independent interest in estimating optimum values. One can see that the problem is indeed more than the sum of its parts, from the fact that we get the optimal bound for linCBwK only when $B \geq \tilde{\Omega}\left(m^{1 / 2} T^{3 / 4}\right)$, unlike either special case for which the optimal bound holds for all $B$ (but is meaningful only for $B=\tilde{\Omega}(m \sqrt{T}))$.
The approach in [3] (for BwK) extends to the case of "static" contexts, ${ }^{4}$ where each arm has a context that doesn't change over time. The OSPP of [4] is not a special case of linCBwK with static contexts; this is one indication of the additional difficulty of dynamic over static contexts.

Other related work. Recently, [17] showed an $O(\sqrt{T})$ regret in the linear contextual setting with a single budget constraint, when costs depend only on contexts and not arms.

Due to space constraints, we have moved many proofs from the main part of the paper to the supplement.

## 2 Preliminaries

### 2.1 Confidence Ellipsoid

Consider a stochastic process which in each round $t$, generates a pair of observations $\left(r_{t}, \boldsymbol{y}_{t}\right)$, such that $r_{t}$ is an unknown linear function of $\boldsymbol{y}_{t}$ plus some 0 -mean bounded noise, i.e., $r_{t}=\boldsymbol{\mu}_{*}^{\top} \boldsymbol{y}_{t}+\eta_{t}$, where $\boldsymbol{y}_{t}, \boldsymbol{\mu}_{*} \in \mathbb{R}^{m},\left|\eta_{t}\right| \leq 2 R$, and $\mathbb{E}\left[\eta_{t} \mid \boldsymbol{y}_{1}, r_{1}, \ldots, \boldsymbol{y}_{t-1}, r_{t-1}, \boldsymbol{y}_{t}\right]=0$.

At any time $t$, a high confidence estimate of the unknown vector $\boldsymbol{\mu}_{*}$ can be obtained by building a "confidence ellipsoid" around the $\ell_{2}$-regularized least-squares estimate $\hat{\boldsymbol{\mu}}_{t}$ constructed from the observations made so far. This technique is common in prior work on linear contextual bandits (e.g., in $[8,11,1])$. For any regularization parameter $\lambda>0$, let

$$
M_{t}:=\lambda I+\sum_{i=1}^{t-1} \boldsymbol{y}_{i} \boldsymbol{y}_{i}^{\top}, \text { and } \hat{\boldsymbol{\mu}}_{t}:=M_{t}^{-1} \sum_{i=1}^{t-1} \boldsymbol{y}_{i} r_{i}
$$

The following result from [1] shows that $\boldsymbol{\mu}_{*}$ lies with high probability in an ellipsoid with center $\hat{\boldsymbol{\mu}}_{t}$. For any positive semi-definite (PSD) matrix $M$, define the $M$-norm as $\|\boldsymbol{\mu}\|_{M}:=\sqrt{\boldsymbol{\mu}^{\top} M \boldsymbol{\mu}}$. The confidence ellipsoid at time $t$ is defined as

$$
C_{t}:=\left\{\boldsymbol{\mu} \in \mathbb{R}^{m}:\left\|\boldsymbol{\mu}-\hat{\boldsymbol{\mu}}_{t}\right\|_{M_{t}} \leq R \sqrt{m \log ((1+t m / \lambda) / \delta)}+\sqrt{\lambda m}\right\}
$$

Lemma 2 (Theorem 2 of [1]). If $\forall t,\left\|\boldsymbol{\mu}_{*}\right\|_{2} \leq \sqrt{m}$ and $\left\|\boldsymbol{y}_{t}\right\|_{2} \leq \sqrt{m}$, then with prob. $1-\delta$, $\boldsymbol{\mu}_{*} \in C_{t}$.

Another useful observation about this construction is stated below. It first appeared as Lemma 11 of [8], and was also proved as Lemma 3 in [11].
Lemma 3 (Lemma 11 of [8]). $\sum_{t=1}^{T}\left\|\boldsymbol{y}_{t}\right\|_{M_{t}^{-1}} \leq \sqrt{m T \log (T)}$.
As a corollary of the above two lemmas, we obtain a bound on the total error in the estimate provided by "any point" from the confidence ellipsoid. (Proof is in Appendix D in the supplement.)

[^2]Corollary 1. For $t=1, \ldots, T$, let $\tilde{\mu}_{t} \in C_{t}$ be a point in the confidence ellipsoid, with $\lambda=1$ and $2 R=1$. Then, with probability $1-\delta$,

$$
\sum_{t=1}^{T}\left|\tilde{\boldsymbol{\mu}}_{t}^{\top} \boldsymbol{y}_{t}-\boldsymbol{\mu}_{*}^{\top} \boldsymbol{y}_{t}\right| \leq 2 m \sqrt{T \log ((1+T m) / \delta) \log (T)} .
$$

### 2.2 Online Learning

Consider a $T$ round game played between an online learner and an adversary, where in round $t$, the learner chooses a $\boldsymbol{\theta}_{t} \in \Omega:=\left\{\boldsymbol{\theta}:\|\boldsymbol{\theta}\|_{1} \leq 1, \boldsymbol{\theta} \geq 0\right\}$, and then observes a linear function $g_{t}: \Omega \rightarrow[-1,1]$ picked by the adversary. The learner's choice $\boldsymbol{\theta}_{t}$ may only depend on learner's and adversary's choices in previous rounds. The goal of the learner is to minimize regret defined as the difference between the learner's objective value and the value of the best single choice in hindsight:

$$
\mathcal{R}(T):=\max _{\boldsymbol{\theta} \in \Omega} \sum_{t=1}^{T} g_{t}(\boldsymbol{\theta})-\sum_{t=1}^{T} g_{t}\left(\boldsymbol{\theta}_{t}\right) .
$$

The multiplicative weight update (MWU) algorithm (generalization by [7]) is a fast and efficient online learning algorithm for this problem. Let $g_{t, j}:=g_{t}\left(\mathbf{1}_{j}\right)$. Then, given a parameter $\epsilon>0$, in round $t+1$, the choice of this algorithm takes the following form,

$$
\boldsymbol{\theta}_{t+1, j}=\frac{w_{t, j}}{1+\sum_{j} w_{t, j}}, \text { where } w_{t, j}= \begin{cases}w_{t-1, j}(1+\epsilon)^{g_{t, j}} & \text { if } g_{t, j}>0  \tag{5}\\ w_{t-1, j}(1-\epsilon)^{-g_{t, j}} & \text { if } g_{t, j} \leq 0\end{cases}
$$

with initialization $w_{0, j}=1$, for all $j=1, \ldots, K$.
Lemma 4. [7] For any $0<\epsilon \leq \frac{1}{2}$, the MWU algorithm provides the following regret bound for the online learning problem described above:

$$
\mathcal{R}(T) \leq \epsilon T+\frac{\log (d+1)}{\epsilon} .
$$

In particular, for $\epsilon=\sqrt{\frac{\log (d+1)}{T}}$, we have $\mathcal{R}(T) \leq \sqrt{\log (d+1) T}$
For the rest of the paper, we refer to the MWU algorithm with $\epsilon=\sqrt{\frac{\log (d+1)}{T}}$ as the online learning (OL) algorithm, and the update in (5) as the OL update at time $t+1$.

## 3 Algorithm

### 3.1 Optimistic estimates of unknown parameters

Let $a_{t}$ denote the arm played by the algorithm at time $t$. In the beginning of every round, we use the outcomes and contexts from previous rounds to construct a confidence ellipsoid for $\boldsymbol{\mu}_{*}$ and every column of $W_{*}$. The construction of confidence ellipsoid for $\boldsymbol{\mu}_{*}$ follows directly from the techniques in Section 2.1 with $y_{t}=\mathbf{x}_{t}\left(a_{t}\right)$ and $r_{t}$ being reward at time $t$. To construct a confidence ellipsoid for a column $j$ of $W_{*}$, we use the techniques in Section 2.1 while substituting $\boldsymbol{y}_{t}=\mathbf{x}_{t}\left(a_{t}\right)$ and $r_{t}=\mathbf{v}_{t}\left(a_{t}\right)_{j}$ for every $j$.
As in Section 2.1, let $M_{t}:=I+\sum_{i=1}^{t-1} \mathbf{x}_{i}\left(a_{i}\right) \mathbf{x}_{i}\left(a_{i}\right)^{\top}$, and construct the regularized least squares estimate for $\boldsymbol{\mu}_{*}, W_{*}$, respectively, as

$$
\begin{align*}
\hat{\boldsymbol{\mu}}_{t} & :=M_{t}^{-1} \sum_{i=1}^{t-1} \mathbf{x}_{i}\left(a_{i}\right) r_{i}\left(a_{i}\right)^{\top}  \tag{6}\\
\hat{W}_{t} & :=M_{t}^{-1} \sum_{i=1}^{t-1} \mathbf{x}_{i}\left(a_{i}\right) \mathbf{v}_{i}\left(a_{i}\right)^{\top} . \tag{7}
\end{align*}
$$

Define confidence ellipsoid for parameter $\boldsymbol{\mu}_{*}$ as

$$
C_{t, 0}:=\left\{\boldsymbol{\mu} \in \mathbb{R}^{m}:\|\boldsymbol{\mu}-\hat{\boldsymbol{\mu}}\|_{M_{t}} \leq \sqrt{m \log ((d+t m d) / \delta)}+\sqrt{m}\right\}
$$

and for every arm $a$, the optimistic estimate of $\boldsymbol{\mu}_{*}$ as:

$$
\begin{equation*}
\tilde{\boldsymbol{\mu}}_{t}(a):=\arg \max _{\boldsymbol{\mu} \in C_{t, 0}} \mathbf{x}_{t}(a)^{\top} \boldsymbol{\mu} . \tag{8}
\end{equation*}
$$

Let $\mathbf{w}_{j}$ denote the $j^{t h}$ column of a matrix $W$. We define a confidence ellipsoid for each column $j$, as

$$
C_{t, j}:=\left\{\mathbf{w} \in \mathbb{R}^{m}:\left\|\mathbf{w}-\hat{\mathbf{w}}_{t j}\right\|_{M_{t}} \leq \sqrt{m \log ((d+t m d) / \delta)}+\sqrt{m}\right\}
$$

and denote by $\mathcal{G}_{t}$, the Cartesian product of all these ellipsoids: $\mathcal{G}_{t}:=\left\{W \in \mathbb{R}^{m \times d}: \mathbf{w}_{j} \in C_{t, j}\right\}$. Note that Lemma 2 implies that $W_{*} \in \mathcal{G}_{t}$ with probability $1-\delta$. Now, given a vector $\boldsymbol{\theta}_{t} \in \mathbb{R}^{d}$, we define the optimistic estimate of the weight matrix at time $t$ w.r.t. $\boldsymbol{\theta}_{t}$, for every $\operatorname{arm} a \in[K]$, as :

$$
\begin{equation*}
\tilde{W}_{t}(a):=\arg \min _{W \in \mathcal{G}_{t}} \mathbf{x}_{t}(a)^{\top} W \boldsymbol{\theta}_{t} . \tag{9}
\end{equation*}
$$

Intuitively, for the reward, we want an upper confidence bound and for the consumption we want a lower confidence bound as an optimistic estimate. This intuition aligns with the above definitions, where the maximizer was used in case of reward and a minimizer was used for consumption. The utility and precise meaning of $\boldsymbol{\theta}_{t}$ will become clearer when we describe the algorithm and present the regret analysis.
Using the definition of $\tilde{\boldsymbol{\mu}}_{t}, \tilde{W}_{t}$, along with the results in Lemma 2 and Corollary 1 about confidence ellipsoids, the following can be derived.
Corollary 2. With probability $1-\delta$, for any sequence of $\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \ldots, \boldsymbol{\theta}_{T}$,

1. $\mathbf{x}_{t}(a)^{\top} \tilde{\boldsymbol{\mu}}_{t}(a) \geq \mathbf{x}_{t}(a)^{\top} \boldsymbol{\mu}_{*}$, for all arms $a \in[K]$, for all time $t$.
2. $\mathbf{x}_{t}(a)^{\top} \tilde{W}_{t}(a) \boldsymbol{\theta}_{t} \leq \mathbf{x}_{t}(a)^{\top} W_{*} \boldsymbol{\theta}_{t}$, for all arms $a \in[K]$, for all time $t$.
3. $\left|\sum_{t=1}^{T}\left(\tilde{\boldsymbol{\mu}}_{t}\left(a_{t}\right)-\boldsymbol{\mu}_{*}\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)\right| \leq(2 m \sqrt{T \log ((1+t m) / \delta) \log (T)})$.
4. $\left\|\sum_{t=1}^{T}\left(\tilde{W}_{t}\left(a_{t}\right)-W_{*}\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)\right\| \leq\left\|\mathbf{1}_{d}\right\|(2 m \sqrt{T \log ((d+\operatorname{tmd}) / \delta) \log (T)})$.

Essentially, the first two claims ensure that we have optimistic estimates, and the last two claims ensure that the estimates quickly converge to the true parameters.

### 3.2 The core algorithm

In this section, we present an algorithm and its analysis, under the assumption that a parameter $Z$ satisfying certain properties is given. Later, we show how to use the first $T_{0}$ rounds to compute such a $Z$, and also bound the additional regret due to these $T_{0}$ rounds. We define $Z$ now.
Assumption 1. Let $Z$ be such that for some universal constants $c, c^{\prime}$, $\frac{O P T}{B} \leq Z \leq c \frac{O P T}{B}+c^{\prime}$.
The algorithm constructs estimates $\hat{\boldsymbol{\mu}}_{t}$ and $\hat{W}_{t}$ as in Section 3.1. It also runs the OL algorithm for an instance of the online learning problem. The vector played by the OL algorithm in time step $t$ is $\boldsymbol{\theta}_{t}$. After observing the context, the optimistic estimates for each arm are then constructed using $\boldsymbol{\theta}_{t}$, as defined in (8) and (9). Intuitively, $\boldsymbol{\theta}_{t}$ is used here as a multiplier to combine different columns of the weight matrix, to get an optimistic weight vector for every arm. An adjusted estimated reward for arm $a$ is then defined by using $Z$ to linearly combine the optimistic estimate of the reward with the optimistic estimate of the consumption, as $\left(\mathbf{x}_{t}(a)^{\top} \tilde{\boldsymbol{\mu}}_{t}(a)\right)-Z\left(\mathbf{x}_{t}(a)^{\top} \tilde{W}_{t}(a) \boldsymbol{\theta}_{t}\right)$. The algorithm chooses the arm which appears to be the best according to the adjusted estimated reward. After observing the resulting reward and consumption vectors, the estimates are updated. The online learning algorithm is advanced by one step, by defining the profit vector to be $\mathbf{v}_{t}\left(a_{t}\right)-\frac{B}{T} \mathbf{1}$. The algorithm ends either after $T$ time steps or as soon as the total consumption exceeds the budget along some dimension.
Theorem 2. Given a $Z$ as per Assumption 1, Algorithm 1 achieves the following, with prob. $1-\delta$ :

$$
\operatorname{regret}(T) \leq O\left(\left(\frac{O P T}{B}+1\right) m \sqrt{T \log (d T / \delta) \log (T)}\right)
$$

(Proof Sketch) We provide a sketch of the proof here, with a full proof given in Appendix E in the supplement. Let $\tau$ be the stopping time of the algorithm. The proof is in 3 steps:
Step 1: Since $\mathbb{E}\left[\mathbf{v}_{t}\left(a_{t}\right) \mid X_{t}, a_{t}, H_{t-1}\right]=W_{*}^{\top} \mathbf{x}_{t}\left(a_{t}\right)$, we apply Azuma-Hoeffding inequality to get that with high probability $\left\|\sum_{t=1}^{\tau} \mathbf{v}_{t}\left(a_{t}\right)-W_{*}^{\top} \mathbf{x}_{t}\left(a_{t}\right)\right\|_{\infty}$ is small. Therefore, we can work with $\sum_{t=1}^{\tau} W_{*}^{\top} \mathbf{x}_{t}\left(a_{t}\right)$ instead of $\sum_{t=1}^{\tau} \mathbf{v}_{t}\left(a_{t}\right)$. A similar application of Azuma-Hoeffding inequality is used to bound the gap $\left|\sum_{t=1}^{\tau} r_{t}\left(a_{t}\right)-\boldsymbol{\mu}_{*}^{\top} \mathbf{x}_{t}\left(a_{t}\right)\right|$, so that a lower bound on $\sum_{t=1}^{\tau} \mu_{*}^{\top} \mathbf{x}_{t}\left(a_{t}\right)$ is sufficient to lower bound the total reward $\sum_{t=1}^{\tau} r_{t}\left(a_{t}\right)$.

```
Algorithm 1 Algorithm for linCBwK, with given \(Z\)
    Initialize \(\boldsymbol{\theta}_{1}\) as per the online learning (OL) algorithm.
    Initialize \(Z\) that satisfies Assumption 1.
    for all \(t=1, \ldots, T\) do
        Observe \(X_{t}\).
        For every \(a \in[K]\), compute \(\tilde{\boldsymbol{\mu}}_{t}(a)\) and \(\tilde{W}_{t}(a)\) as per (8) and (9) respectively.
        Play the \(\operatorname{arm} a_{t}:=\arg \max _{a \in[K]} \mathbf{x}_{t}(a)^{\top}\left(\tilde{\boldsymbol{\mu}}_{t}(a)-Z \tilde{W}_{t}(a) \boldsymbol{\theta}_{t}\right)\).
        Observe \(r_{t}\left(a_{t}\right)\) and \(\mathbf{v}_{t}\left(a_{t}\right)\).
        If for some \(j=1 . . d, \sum_{t^{\prime}<t} \mathbf{v}_{t^{\prime}}\left(a_{t^{\prime}}\right) \cdot \mathbf{e}_{j} \geq B\) then EXIT.
        Use \(\mathbf{x}_{t}\left(a_{t}\right), r_{t}\left(a_{t}\right)\) and \(\mathbf{v}_{t}\left(a_{t}\right)\) to obtain \(\hat{\boldsymbol{\mu}}_{t+1}, \hat{W}_{t+1}\) and \(\mathcal{G}_{t+1}\).
        Choose \(\boldsymbol{\theta}_{t+1}\) using the OL update (refer to (5)) with \(g_{t}\left(\boldsymbol{\theta}_{t}\right):=\boldsymbol{\theta}_{t} \cdot\left(\mathbf{v}_{t}\left(a_{t}\right)-\frac{B}{T} \mathbf{1}\right)\).
    end for
```

Step 2: Using Corollary 2, with high probability, we can bound $\left\|\sum_{t=1}^{T}\left(W_{*}-\tilde{W}_{t}\left(a_{t}\right)\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)\right\|_{\infty}$. It is therefore sufficient to work with the sum of vectors $\tilde{W}_{t}\left(a_{t}\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)$ instead of $W_{*}^{\top} \mathbf{x}_{t}\left(a_{t}\right)$, and similarly with $\tilde{\boldsymbol{\mu}}_{t}\left(a_{t}\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)$ instead of $\mu_{*}^{\top} \mathbf{x}_{t}\left(a_{t}\right)$.
Step 3: The proof is completed by showing the desired bound on OPT $-\sum_{t=1}^{\tau} \tilde{\boldsymbol{\mu}}_{t}\left(a_{t}\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)$. This part is similar to the online stochastic packing problem; if the actual reward and consumption vectors were $\tilde{\boldsymbol{\mu}}_{t}\left(a_{t}\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)$ and $\tilde{W}_{t}\left(a_{t}\right)^{\top} \mathbf{x}_{t}\left(a_{t}\right)$, then it would be exactly that problem. We adapt techniques from [4]: use the OL algorithm and the $Z$ parameter to combine constraints into the objective. If a dimension is being consumed too fast, then the multiplier for that dimension should increase, making the algorithm to pick arms that are not likely to consume too much along this dimension. Regret is then bounded by a combination of the online learning regret and the error in the optimistic estimates.

### 3.3 Algorithm with $Z$ computation

In this section, we present a modification of Algorithm 1 which computes the required parameter $Z$ that satisfies Assumption 1, and therefore does not need to be provided with a $Z$ as input. The algorithm computes $Z$ using observations from the first $T_{0}$ rounds. Once $Z$ is computed, Algorithm 1 can be run for the remaining time steps. However, it needs to be modified slightly to take into account the budget consumed during the first $T_{0}$ rounds. We handle this by using a smaller budget $B^{\prime}=B-T_{0}$ in the computations for the remaining rounds. The modified algorithm is given below.

```
Algorithm 2 Algorithm for linCBwK, with \(Z\) computation
    Inputs: \(B, T_{0}, B^{\prime}=B-T_{0}\)
    Using observations from first \(T_{0}\) rounds, compute a \(Z\) that satisfies Assumption 1.
    Run Algorithm 1 for \(T-T_{0}\) rounds and budget \(B^{\prime}\).
```

Next, we provide the details of how to compute $Z$ from observations in the first $T_{0}$ rounds, and how to choose $T_{0}$. We provide a method that takes advantage of the linear structure of the problem, and explores in the $m$-dimensional space of contexts and weight vectors to obtain bounds independent of $K$. In every round $t=1, \ldots, T_{0}$, after observing $X_{t}$, let $p_{t} \in \Delta^{[K]}$ be

$$
\begin{align*}
p_{t} & :=\arg \max _{p \in \Delta^{[K]}}\left\|X_{t} p\right\|_{M_{t}^{-1}}  \tag{10}\\
\text { where } M_{t} & :=I+\sum_{i=1}^{t-1}\left(X_{i} p_{i}\right)\left(X_{i} p_{i}\right)^{\top} . \tag{11}
\end{align*}
$$

Select arm $a_{t}=a$ with probability $p_{t}(a)$. In fact, since $M_{t}$ is a PSD matrix, due to convexity of the function $\left\|X_{t} p\right\|_{M_{t}^{-1}}^{2}$, it is the same as playing $a_{t}=\arg \max _{a \in[K]}\left\|\mathbf{x}_{t}(a)\right\|_{M_{t}^{-1}}$. Construct estimates $\hat{\boldsymbol{\mu}}, \hat{W}_{t}$ of $\boldsymbol{\mu}_{*}, W_{*}$ at time $t$ as

$$
\hat{\mu}_{t}:=M_{t}^{-1} \sum_{i=1}^{t-1}\left(X_{i} p_{i}\right) r_{i}\left(a_{i}\right), \quad \hat{W}_{t}:=M_{t}^{-1} \sum_{i=1}^{t-1}\left(X_{i} p_{i}\right) \mathbf{v}_{i}\left(a_{i}\right)^{\top} .
$$

And, for some value of $\gamma$ defined later, obtain an estimate $\mathrm{OPT}^{\gamma}$ of OPT as:

$$
\text { OPT }^{\gamma}:=\quad \begin{array}{rc}
\max _{\pi} & \frac{T}{T_{0}} \sum_{i=1}^{T_{0}} \hat{\boldsymbol{\mu}}_{i}^{\top} X_{i} \pi\left(X_{i}\right)  \tag{12}\\
\text { such that } & \frac{T}{T_{0}} \sum_{i=1}^{T_{0}} \hat{W}_{i}^{\top} X_{i} \pi\left(X_{i}\right) \leq B+\gamma .
\end{array}
$$

For an intuition about the choice of arm in (10), observe from the discussion in Section 2.1 that every column $\mathbf{w}_{* j}$ of $W_{*}$ is guaranteed to lie inside the confidence ellipsoid centered at column $\hat{\mathbf{w}}_{t j}$ of $\hat{W}_{t}$, namely the ellipsoid, $\left\|\mathbf{w}-\hat{\mathbf{w}}_{t j}\right\|_{M_{t}}^{2} \leq 4 m \log (T m / \delta)$. Note that this ellipsoid has principle axes as eigenvectors of $M_{t}$, and the length of the semi-principle axes is given by the inverse eigenvalues of $M_{t}$. Therefore, by maximizing $\left\|X_{t} p\right\|_{M_{t}^{-1}}$ we are choosing the context closest to the direction of the longest principal axis of the confidence ellipsoid, i.e. in the direction of the maximum uncertainty. Intuitively, this corresponds to pure exploration: by making an observation in the direction where uncertainty is large we can reduce the uncertainty in our estimate most effectively.
A more algebraic explanation is as follows. In order to get a good estimate of OPT by OP̂T ${ }^{\gamma}$, we want the estimates $\hat{W}_{t}$ and $W_{*}$ (and, $\hat{\boldsymbol{\mu}}$ and $\boldsymbol{\mu}_{*}$ ) to be close enough so that $\left\|\sum_{t=1}^{T_{0}}\left(\hat{W}_{t}-\hat{W}_{*}\right)^{\top} X_{t} \pi\left(X_{t}\right)\right\|_{\infty}$ (and, $\left.\left|\sum_{t=1}^{T_{0}}\left(\hat{\boldsymbol{\mu}}_{t}-\boldsymbol{\mu}_{*}\right)^{\top} X_{t} \pi\left(X_{t}\right)\right|\right)$ is small for all policies $\pi$, and in particular for sample optimal policies. Now, using Cauchy-Schwartz these are bounded by

$$
\begin{gathered}
\left.\sum_{t=1}^{T_{0}}\left\|\hat{\boldsymbol{\mu}}_{t}-\boldsymbol{\mu}_{*}\right\|_{M_{t}} \| X_{t} \pi\left(X_{t}\right)\right) \|_{M_{t}^{-1}}, \text { and } \\
\sum_{t=1}^{T_{0}}\left\|\hat{W}_{t}-W_{*}\right\|_{M_{t}}\left\|X_{t} \pi\left(X_{t}\right)\right\|_{M_{t}^{-1}}
\end{gathered}
$$

where we define $\|W\|_{M}$, the $M$-norm of matrix $W$ to be the max of column-wise $M$-norms. Using Lemma 2, the term $\left\|\hat{\boldsymbol{\mu}}_{t}-\boldsymbol{\mu}_{*}\right\|_{M_{t}}$ is bounded by $2 \sqrt{m \log \left(T_{0} m / \delta\right)}$, and $\left\|\hat{W}_{t}-W_{*}\right\|_{M_{t}}$ is bounded by $2 \sqrt{m \log \left(T_{0} m d / \delta\right)}$, with probability $1-\delta$. Lemma 3 bounds the second term $\sum_{t=1}^{T_{0}}\left\|X_{t} \pi\left(X_{t}\right)\right\|_{M_{t}^{-1}}$ but only when $\pi$ is the played policy. This is where we use that the played policy $p_{t}$ was chosen to maximize $\left\|X_{t} p_{t}\right\|_{M_{t}^{-1}}$, so that $\sum_{t=1}^{T_{0}}\left\|X_{t} \pi\left(X_{t}\right)\right\|_{M_{t}^{-1}} \leq \sum_{t=1}^{T_{0}}\left\|X_{t} p_{t}\right\|_{M_{t}^{-1}}$ and the bound $\sum_{t=1}^{T_{0}}\left\|X_{t} p_{t}\right\|_{M_{t}^{-1}} \leq \sqrt{m T_{0} \log \left(T_{0}\right)}$ given by Lemma 3 actually bounds $\sum_{t=1}^{T_{0}}\left\|X_{t} \pi\left(X_{t}\right)\right\|_{M_{t}^{-1}}$ for all $\pi$. Combining, we get a bound of $2 m \sqrt{T_{0} \log \left(T_{0}\right) \log \left(T_{0} d / \delta\right)}$ on deviations $\| \sum_{t=1}^{T_{0}}\left(\hat{W}_{t}-\right.$ $\left.\hat{W}_{*}\right)^{\top} X_{t} \pi\left(X_{t}\right) \|_{\infty}$ and $\left|\sum_{t=1}^{T_{0}}\left(\hat{\boldsymbol{\mu}}_{t}-\boldsymbol{\mu}_{*}\right)^{\top} X_{t} \pi\left(X_{t}\right)\right|$ for all $\pi$.
We prove the following lemma.
Lemma 5. For $\gamma=\left(\frac{T}{T_{0}}\right) 2 m \sqrt{T_{0} \log \left(T_{0}\right) \log \left(T_{0} d / \delta\right)}$, with probability $1-O(\delta)$,

$$
O P T-2 \gamma \leq O \hat{P} T^{2 \gamma} \leq O P T+9 \gamma\left(\frac{O P T}{B}+1\right)
$$

Corollary 3. Set $Z=\frac{\left(O \hat{P} T^{2 \gamma}+2 \gamma\right)}{B}+1$, with the above value of $\gamma$. Then, with probability $1-O(\delta)$,

$$
\frac{O P T}{B}+1 \leq Z \leq\left(1+\frac{11 \gamma}{B}\right)\left(\frac{O P T}{B}+1\right) .
$$

Corollary 3 implies that as long as $B \geq \gamma$, i.e., $B \geq \tilde{\Omega}\left(\frac{m T}{\sqrt{T_{0}}}\right), Z$ is a constant factor approximation of $\frac{\mathrm{OPT}}{B}+1 \geq Z^{*}$, therefore Theorem 2 should provide an $\tilde{O}\left(\left(\frac{\mathrm{OPT}}{B}+1\right) m \sqrt{T}\right)$ regret bound. However, this bound does not account for the budget consumed in the first $T_{0}$ rounds. Considering that (at most) $T_{0}$ amount can be consumed from the budget in the first $T_{0}$ rounds, we have an additional regret of $\frac{\text { OPT }}{B} T_{0}$. Further, since we have $B^{\prime}=B-T_{0}$ budget for remaining $T-T_{0}$ rounds, we need a $Z$ that satisfies the required assumption for $B^{\prime}$ instead of $B$ (i.e., we need $\frac{\mathrm{OPT}}{B^{\prime}} \leq Z \leq O(1)\left(\frac{\mathrm{OPT}}{B^{\prime}}+1\right)$ ). If $B \geq 2 T_{0}$, then, $B^{\prime} \geq B / 2$, and using 2 times the $Z$ computed in Corollary 3 would satisfy the required assumption.
Together, these observations give Theorem 3.
Theorem 3. Using Algorithm 2 with $T_{0}$ such that $B>\max \left\{2 T_{0}, m T / \sqrt{T_{0}}\right\}$, and twice the $Z$ given by Corollary 3, we get a high probability regret bound of

$$
\tilde{O}\left(\left(\frac{O P T}{B}+1\right)\left(T_{0}+m \sqrt{T}\right)\right)
$$

In particular, for $B>m^{1 / 2} T^{3 / 4}$ and $m \leq \sqrt{T}$, we can use $T_{0}=m \sqrt{T}$ to get a regret bound of

$$
\tilde{O}\left(\left(\frac{O P T}{B}+1\right) m \sqrt{T}\right)
$$

## References

[1] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In NIPS, 2012.
[2] A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, and R. E. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In ICML 2014, June 2014.
[3] S. Agrawal and N. R. Devanur. Bandits with concave rewards and convex knapsacks. In Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC '14, 2014.
[4] S. Agrawal and N. R. Devanur. Fast algorithms for online stochastic convex programming. In SODA, pages 1405-1424, 2015.
[5] S. Agrawal, N. R. Devanur, and L. Li. An efficient algorithm for contextual bandits with knapsacks, and an extension to concave objectives. In COLT, 2016.
[6] S. Agrawal, Z. Wang, and Y. Ye. A dynamic near-optimal algorithm for online linear programming. Operations Research, 62:876-890, 2014.
[7] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. Theory of Computing, 8(6):121-164, 2012.
[8] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. J. Mach. Learn. Res., 3, Mar. 2003.
[9] A. Badanidiyuru, R. Kleinberg, and A. Slivkins. Bandits with knapsacks. In FOCS, pages 207-216, 2013.
[10] A. Badanidiyuru, J. Langford, and A. Slivkins. Resourceful contextual bandits. In Proceedings of The Twenty-Seventh Conference on Learning Theory (COLT-14), pages 1109-1134, 2014.
[11] W. Chu, L. Li, L. Reyzin, and R. E. Schapire. Contextual Bandits with Linear Payoff Functions. In AISTATS, 2011.
[12] V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic Linear Optimization under Bandit Feedback. In COLT, 2008.
[13] N. R. Devanur and T. P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In EC, 2009.
[14] N. R. Devanur, K. Jain, B. Sivan, and C. A. Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In EC, 2011.
[15] J. Feldman, M. Henzinger, N. Korula, V. S. Mirrokni, and C. Stein. Online stochastic packing applied to display ad allocation. In Proceedings of the 18th Annual European Conference on Algorithms: Part I, ESA'10, 2010.
[16] V. Guruswami and P. Raghavendra. Hardness of learning halfspaces with noise. SIAM Journal on Computing, 39(2):742-765, 2009.
[17] H. Wu, R. Srikant, X. Liu, and C. Jiang. Algorithms with logarithmic or sublinear regret for constrained contextual bandits. In Proceedings of the 28th International Conference on Neural Information Processing Systems (NIPS), 2015.


[^0]:    *Columbia University. sa3305@columbia.edu.
    ${ }^{\dagger}$ Microsoft Research. nikdev@microsoft. com.

[^1]:    ${ }^{3}$ Similar to the regret bounds for linear contextual bandits [8, 1, 11].

[^2]:    ${ }^{4}$ It was incorrectly claimed in [3] that the approach can be extended to dynamic contexts without much modifications.

