

# Online Learning with Vector Costs and Bandits with Knapsacks

**Thomas Kesselheim**

*Institute of Computer Science, University of Bonn.*

THOMAS.KESSELHEIM@UNI-BONN.DE

**Sahil Singla**

*Department of Computer Science at Princeton University and School of Mathematics at Institute for Advanced Study. Supported in part by the Schmidt Foundation.*

SINGLA@CS.PRINCETON.EDU

**Editors:** Jacob Abernethy and Shivani Agarwal

## Abstract

We introduce online learning with vector costs (OLVC<sub>p</sub>) where in each time step  $t \in \{1, \dots, T\}$ , we need to play an action  $i \in \{1, \dots, n\}$  that incurs an unknown vector cost in  $[0, 1]^d$ . The goal of the online algorithm is to minimize the  $\ell_p$  norm of the sum of its cost vectors. This captures the classical online learning setting for  $d = 1$ , and is interesting for general  $d$  because of applications like online scheduling where we want to balance the load between different machines (dimensions).

We study OLVC<sub>p</sub> in both stochastic and adversarial arrival settings, and give a general procedure to reduce the problem from  $d$  dimensions to a single dimension. This allows us to use classical online learning algorithms in both full and bandit feedback models to obtain (near) optimal results. In particular, we obtain a single algorithm (up to the choice of learning rate) that gives sublinear regret for stochastic arrivals and a tight  $O(\min\{p, \log d\})$  competitive ratio for adversarial arrivals.

The OLVC<sub>p</sub> problem also occurs as a natural subproblem when trying to solve the popular Bandits with Knapsacks (BWK) problem. This connection allows us to use our OLVC<sub>p</sub> techniques to obtain (near) optimal results for BWK in both stochastic and adversarial settings. In particular, we obtain a tight  $O(\log d \cdot \log T)$  competitive ratio algorithm for adversarial BWK, which improves over the  $O(d \cdot \log T)$  competitive ratio algorithm of [Immorlica et al. \(2019\)](#).

## 1. Introduction

The field of online learning has seen tremendous progress in the last three decades since the seminal work of [Littlestone and Warmuth \(1989\)](#). This is primarily because it is the simplest model to study exploration vs. exploitation trade-off while still being powerful enough to model several applications like clinical trials, ad-allocation, and portfolio optimization. The classic setup of online learning is that we are given a set  $\mathcal{A}$  of  $n$  actions and in each time step  $t \in [T]$  we take an action  $x^{(t)} \in \Delta_n$ <sup>1</sup>. After taking the action we incur a *scalar* cost  $c^{(t)} x^{(t)}$ , where  $c^{(t)} \in [0, 1]^{1 \times n}$  is a row vector of costs, and receive a *feedback* about  $c^{(t)}$ . The goal of the algorithm is to minimize its total cost compared to that of the best *fixed action*  $x^* \in \Delta_n$  chosen in hindsight. Several variants of this basic problem have been studied in the literature, depending on the amount of feedback that the algorithm receives (full vs. partial) and who draws the costs  $c^{(t)}$  (adversary vs. stochastic) (see books [Cesa-Bianchi and Lugosi \(2006\)](#); [Bubeck and Cesa-Bianchi \(2012\)](#); [Hazan \(2016\)](#); [Slivkins \(2019\)](#); [Lattimore and Szepesvári \(2020\)](#)).

1. We use the notation  $[T] := \{1, 2, \dots, T\}$  and write  $\Delta_n$  to mean the simplex  $\{\mathbf{x} \in [0, 1]^n \mid \sum_{i=1}^n x_i = 1\}$ .

In this work we study online learning problems where the actions incur *vector costs* in  $[0, 1]^d$ . The algorithm wants to minimize the  $\ell_p$  norm, given some integer  $p \geq 1$ , of the total cost vector. Formally, in the *online learning with vector costs* (OLVC $_p$ ) problem there is a set  $\mathcal{A}$  of  $n$  actions and in each time step  $t$  we take an action  $x^{(t)} \in \Delta_n$ . After taking the action we incur a *vector cost*  $C^{(t)}x^{(t)}$ , where  $C^{(t)} \in [0, 1]^{d \times n}$  is a cost matrix, and receive a feedback about  $C^{(t)}$ . The goal of the algorithm is to minimize the  $\ell_p$  norms of its total cost vector compared to that of a fixed benchmark distribution  $x^* \in \Delta_n$  over the actions<sup>2</sup>. This problem clearly captures classic online learning for  $d = 1$ . Again we want to study the problem depending on both the amount of feedback and on who draws the cost matrices.

The motivation to study online learning with vector costs is twofold. We first argue below that it is a natural problem in itself because it models the classical scheduling problem of online load balancing. Next we argue that the problem is also interesting because it is a natural step towards designing (optimal) algorithms for the Bandits with Knapsacks (BWK) problem, both in the stochastic and adversarial settings. The BWK problem has been applications like dynamic item pricing, repeated auctions, and dynamic procurement (e.g., see [Badanidiyuru et al. \(2018\)](#); [Agrawal and Devanur \(2014\)](#); [Sankararaman and Slivkins \(2018\)](#); [Immorlica et al. \(2019\)](#)).

### 1.1. Online Load Balancing

In the classical *online generalized load balancing* (OGLB $_p$ ) problem, a sequence of  $T$  jobs arrive one-by-one and the algorithm has to process them on  $d$  machines to minimize the  $\ell_p$  norm of the total load vector. Each job can be processed in  $n$  ways, and playing  $x^{(t)} \in \Delta_n$  incurs a vector load  $C^{(t)}x^{(t)}$ , where  $C^{(t)} \in [0, 1]^{d \times n}$  is a load matrix. It generalizes the fundamental makespan minimization problem where we want to minimize the maximum load, i.e., the  $\ell_\infty$  norm. The motivation to study  $\ell_p$  norms for arbitrary  $p$  comes from noticing that  $p = 1$  captures the total load,  $p = \infty$  captures the maximum load, and  $p$  between 1 and  $\infty$  interpolates between the two extremes. The problem and its special cases have been well studied both for adversarial loads ([Awerbuch et al. \(1995\)](#); [Caragiannis \(2008\)](#); [Im et al. \(2019\)](#)) and stochastic loads ([Molinaro \(2017\)](#)).

The major difference between the classical OGLB $_p$  and our OLVC $_p$  problems is that in the former the algorithm gets to see the load/cost matrix  $C^{(t)}$  *before* playing the action, whereas in the latter the algorithm receives a full/partial feedback only *after* playing the action. Depending on the application in hand, both models make sense as it may or may not be possible to find the cost matrix before playing the action. Of course, this makes a huge difference in what guarantees one can hope to achieve. Indeed, in OGLB $_p$  the benchmark is an arbitrary sequence of actions (policy) chosen in hindsight, while in OLVC $_p$  the benchmark is a fixed distribution  $x^*$  of actions. Given a benchmark, we can compare different algorithms using the *regret* or the *competitive-ratio* framework (defined below), depending on whether the ratio of the costs of the algorithm and the benchmark tend towards 1, or not, as  $T \rightarrow \infty$ .

In OLVC $_p$  the benchmark is a fixed distribution  $x^* \in \Delta_n$  on actions. Its cost is defined to be

$$\text{OPT}_{\text{OLVC}} := \left\| \sum_{t=1}^T C^{(t)} \cdot x^* \right\|_p \quad \text{or} \quad \text{OPT}_{\text{OLVC}} := \mathbb{E} \left[ \left\| \sum_{t=1}^T C^{(t)} \cdot x^* \right\|_p \right], \quad (1)$$

2. The optimal fixed  $x^*$  is a distribution over actions, rather than being a fixed deterministic action as in the  $d = 1$  case. E.g., suppose  $d = 2$  and  $p = \infty$  and there are two actions with the cost matrix  $C^{(t)}$  always being a 2-dimensional identity matrix. Here the optimal distribution is  $(0.5, 0.5)$ , which is a constant factor better than any fixed action.

depending on the adversarial or the stochastic setting, respectively. The total cost of the algorithm is  $\text{ALG}_{\text{OLVC}} := \mathbb{E}[\|\sum_{t=1}^T C^{(t)} \cdot x^{(t)}\|_p]$ , where the expectation is taken over any randomization of the algorithm and of the cost matrices. We say an algorithm is  $\alpha$ -competitive if  $\text{ALG}_{\text{OLVC}} \leq \alpha \cdot \text{OPT}_{\text{OLVC}} + o(T)$ . For 1-competitive algorithms we define *regret* to be  $\text{ALG}_{\text{OLVC}} - \text{OPT}_{\text{OLVC}}$ .

The following is our main result for  $\text{OLVC}_p$ , which is formally stated in Theorems 1 and 2.

**Main Result 1.1** *For stochastic  $\text{OLVC}_p$  there exists an algorithm that guarantees  $\text{ALG}_{\text{OLVC}} \leq \text{OPT}_{\text{OLVC}} + o(T)$ . For adversarial  $\text{OLVC}_p$  there exists an algorithm that guarantees  $\text{ALG}_{\text{OLVC}} \leq O(\min\{p, \log d\}) \cdot \text{OPT}_{\text{OLVC}} + o(T)$ .*

So, for  $\ell_\infty$  norm, we get an  $O(\log d)$ -competitive algorithm for OLVC with adversarial arrivals. A powerful feature of this result is that the same algorithm (up to the choice of the learning-rate parameter) achieves (near) optimal guarantees in both the stochastic and adversarial settings.

This work is not the first one to consider online learning problems with vector costs. Indeed, Agrawal and Devanur (2014) study the *Bandits with Convex Knapsacks* problem where one also incurs vector costs, and the regret is measured in terms of the distance from a given convex body. Since  $\ell_p$  norm is convex, their UCB-style algorithms can be used to obtain sublinear regrets for our *stochastic*  $\text{OLVC}_p$ . We are however, to the best of our knowledge, the first ones to consider  $\text{OLVC}_p$  for adversarial arrivals. Our challenges and techniques are very different from Agrawal and Devanur (2014) for two reasons. On the one hand, UCB-style techniques do not apply for adversarial arrivals. On the other hand, we show in §B that every online algorithm is  $\Omega(\min\{p, \log d\})$ -competitive for adversarial  $\text{OLVC}_p$ . The fact that our adversarial  $\text{OLVC}_p$  algorithm also gives near-optimal guarantees for stochastic  $\text{OLVC}_p$  is a bonus.

## 1.2. Bandits with Knapsacks

The Bandits with Knapsacks (BWK) problem was introduced in Badanidiyuru et al. (2018) to handle applications of multi-armed bandits with resource constraints. It is very similar to the above OLVC problem, but is in a maximization setting. We define its generalization to arbitrary  $\ell_p$  norms, denoted by  $\text{BWK}_p$  as follows (for  $p = \infty$  this is the same as BWK). We are given a budget  $B \geq 0$  and a set  $\mathcal{A}$  of  $n$  actions. In each time step  $t \in [T]$ , we take an action  $x^{(t)} \in \Delta_n$ . After taking the action, we receive a *scalar* reward  $r^{(t)}x^{(t)}$  and incur a *vector* cost  $C^{(t)}x^{(t)}$ , where  $r^{(t)} \in [0, 1]^{1 \times n}$  is a row vector of rewards and  $C^{(t)} \in [0, 1]^{d \times n}$  is a cost matrix. Moreover, we also receive a partial/full feedback about  $C^{(t)}$  and  $r^{(t)}$ . We assume there is a *null action* in  $\mathcal{A}$  that gives 0 reward and  $\mathbf{0}_d$  vector cost<sup>3</sup>, which allows us to skip some time steps. Our goal is to maximize the total reward received while ensuring that the  $\ell_p$  norm of the total cost vector is less than  $B$ . After exhausting the budget, we are only allowed to choose the null action (thus we obtain no further reward).

The benchmark for  $\text{BWK}_p$  is any fixed distribution  $x^* \in \Delta_n$  over actions. Let  $\tau^*$  be the time step at which the benchmark exhausts its budget, i.e.,  $\tau^* \leq T$  is maximal such that  $\|\sum_{t=1}^{\tau^*} C^{(t)}x^*\|_p \leq B$  or  $\|\mathbb{E}[\sum_{t=1}^{\tau^*} C^{(t)}x^*]\|_p \leq B$  depending on adversarial/stochastic setting. Now the optimal reward  $\text{OPT}_{\text{BWK}}$  is defined to be  $\sum_{t=1}^{\tau^*} r^{(t)}x^*$  or  $\mathbb{E}[\sum_{t=1}^{\tau^*} r^{(t)}x^*]$ , respectively, and the value of the algorithm  $\text{ALG}_{\text{BWK}} := \mathbb{E}[\sum_{t=1}^T r^{(t)}x^{(t)}]$ , where the expectation is taken over any randomization of the algorithm and the cost matrices. To ensure that the competitive ratio  $\alpha$  is  $\geq 1$  and the regret is

3. We use the notation  $\mathbf{0}_d$  (or  $\mathbf{1}_d$ ) to mean a  $d$ -dimensional vector of all 0s (or 1s).

non-negative for  $\text{BWK}_p$ , we say an algorithm is  $\alpha$ -competitive if  $\text{ALG}_{\text{OLVC}} \geq \frac{1}{\alpha} \cdot \text{OPT}_{\text{OLVC}} - o(T)$ . For 1-competitive algorithms, we define its *regret* to be  $\text{OPT}_{\text{BWK}} - \text{ALG}_{\text{BWK}}$ .

The BWK problem has been generalized in various directions since the original work of [Badanidiyuru et al. \(2018\)](#) (see related work in §1.3). All these works were, however, only in the stochastic setting. In a very recent work, [Immorlica et al. \(2019\)](#) initiated the study of BWK in an adversarial setting where both the costs and the rewards are chosen by an adversary. Since sub-linear regrets are impossible here, they design algorithms with a small competitive ratio  $\alpha$ . In particular, assuming the value of  $\text{OPT}_{\text{BWK}}$  is known, they give an  $O(d)$ -competitive algorithm and show that every online algorithm is  $\Omega(1)$ -competitive. Without this assumption, they give an  $O(d \cdot \log T)$ -competitive algorithm and show that every online algorithm is  $\Omega(\log T)$ -competitive.

Although [Immorlica et al. \(2019\)](#) give the optimal dependency on the time-horizon  $T$  in the competitive ratio of adversarial BWK, they leave open what is the optimal dependency on the number of dimensions  $d$ . Understanding the optimal dependency on  $d$  is useful for applications with large  $d$ , e.g., in online ad-allocations every bidder has a budget constraint and  $d$  corresponds to the number of bidders ([Mehta, 2012](#)). Our second main result (formally proven in Theorems 5 and 6) resolves this open question of [Immorlica et al. \(2019\)](#) and gives the optimal competitive ratio for adversarial BWK, and its generalization to adversarial  $\text{BWK}_p$ .

**Main Result 1.2** *For stochastic  $\text{BWK}_p$  there exists an algorithm that guarantees  $\text{ALG}_{\text{BWK}} \geq \text{OPT}_{\text{BWK}} - o(T)$ . For adversarial  $\text{BWK}_p$  there exists an algorithm that guarantees  $\text{ALG}_{\text{BWK}} \geq \frac{1}{O(\min\{p, \log d\})} \cdot \text{OPT}_{\text{BWK}} - o(T)$  assuming it is given the value  $\text{OPT}_{\text{BWK}}$ , and  $\text{ALG}_{\text{BWK}} \geq \frac{1}{O(\min\{p, \log d\} \log T)} \cdot \text{OPT}_{\text{BWK}} - o(T)$  when  $\text{OPT}_{\text{BWK}}$  is not known.*

Note that an immediate corollary of the result is an  $O(\log d)$ -competitive algorithm for adversarial BWK when  $\text{OPT}_{\text{BWK}}$  is known, and an  $O(\log d \cdot \log T)$ -competitive algorithm otherwise. This improves on the  $O(d)$ -competitive and the  $O(d \cdot \log T)$ -competitive algorithms of [Immorlica et al. \(2019\)](#) in the respective settings. Our proof crucially relies on the connections between  $\text{BWK}_p$  and  $\text{OLVC}_p$ , after we Lagrangify the budget constraint of  $\text{BWK}_p$ . In §B we also show that our competitive ratios for adversarial arrivals in Main Result 1.2 are optimal.

### 1.3. Further Related Work

Since the field of online learning is vast, we only discuss the most relevant work and refer the reader to the books mentioned in the introduction for other references. At a high-level, our proofs go by reducing the  $d$ -dimensional problem to a surrogate classical 1-dimensional online learning problem against an *adaptive* adversary. Now we run classical algorithms like Hedge ([Freund and Schapire, 1999](#)) in the full feedback setting and EXP-3.P ([Auer et al., 2002](#)) in the bandit feedback setting.

Since the original work of [Badanidiyuru et al. \(2018\)](#) (conference version appeared in 2013), the BWK problem has been generalized in various directions like concave-rewards with convex-costs ([Agrawal and Devanur, 2014](#)), contextual bandits ([Agrawal and Devanur, 2016](#); [Agrawal et al., 2016](#)), and combinatorial semi-bandits ([Sankararaman and Slivkins, 2018](#)). All these works are in the stochastic setting and use UCB-style algorithms to obtain sub-linear regrets. [Immorlica et al. \(2019\)](#) introduce the adversarial BWK problem and show that their techniques also apply to some of these directions for adversarial arrivals. Another relevant work is that of [Rangi et al. \(2019\)](#) who consider adversarial BWK but only for  $d = 1$ .

The competitive analysis framework has been popular in the online algorithms community since at least the work of [Sleator and Tarjan \(1985\)](#). We refer the readers to the books ([Borodin and El-Yaniv, 2005](#); [Buchbinder et al., 2009](#)). Unlike online learning, these results are in a setting where the algorithm gets full-information about the next arrival before making its decision. In recent years, competitive analysis has also become popular for online learning. In particular, [Kakade et al. \(2009\)](#) showed that for any linear combinatorial optimization problem for which we know an  $\alpha$ -approximation offline algorithm, we can also obtain an  $\alpha$ -competitive online combinatorial optimization algorithm against the benchmark that always plays a fixed feasible solution chosen in hindsight. Their result does not apply to our setting as our objective is not additive over time steps.

Finally, we remark that vector losses also appear in the classical Blackwell’s approachability theorem ([Blackwell et al., 1956](#)). Due to the work of [Abernethy et al. \(2011\)](#), it is known this is equivalent to no-regret learning. [Agrawal and Devanur \(2014\)](#) exploit this connection to design a fast algorithm for stochastic BwK. These techniques, however, seem less relevant for our adversarial arrival problems where sublinear regrets are impossible.

## 2. Our Approach via a One-Dimensional Surrogate Online Learning Problem

The idea of our algorithms is to reduce the multi-dimensional problem with vector costs to a single dimensional problem. In each step  $t$ , each action  $i$  is assigned a scalar loss  $c_i^{(t)} \in [0, \|\mathbf{1}\|_p]$ . This loss will mirror the increase of the  $\ell_p$  norm that would be incurred by choosing action  $i$  in step  $t$ . In this surrogate problem, standard online learning algorithms can be applied.

In more detail, for any time  $t \in \{0, \dots, T\}$  and any action sequence  $x^{(1)}, \dots, x^{(t-1)}$  of the algorithm, let  $\mathbf{\Lambda}^{(t)} := \sum_{s=1}^t C^{(s)} \cdot x^{(s)}$  denote the load vector after step  $t$ . That is, in the load balancing setting, our goal is to minimize  $\|\mathbf{\Lambda}^{(T)}\|_p$ .

To define the cost of an action, we use a smooth approximation  $\Psi(\cdot)$  of the  $\ell_p$  norm, which we introduce below. We set the cost of action/expert  $i \in [n]$  in step  $t$  in the surrogate problem to be

$$c_i^{(t)} := \langle C^{(t)} e_i, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle,$$

where  $e_i$  is the  $i$ -th unit vector. The idea is to choose  $\Psi$  so that  $\sum_{t=1}^T \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \approx \|\mathbf{\Lambda}^{(T)}\|_p$ . Note that  $c^{(t)}$  depends on  $x^{(1)}, \dots, x^{(t-1)}$  via  $\mathbf{\Lambda}^{(t-1)}$ . So, the surrogate problem is against an *adaptive* (non-oblivious) adversary.

**A Smooth Approximation of the Norm.** Let us first see an example why smoothing the  $\ell_p$  norm is important. Consider  $p = \infty$  and there are only two actions: the first action puts  $(1 - \epsilon)$  load on every dimension for  $\epsilon \rightarrow 0$  and the second action puts 1 load on one of the  $d$  dimensions chosen at random. Here a greedy algorithm w.r.t. the  $\ell_\infty$  norm will chose the first action (to save an  $\epsilon$ ), but we should have chosen the second action. Hence the greedy algorithm is  $\Omega(d)$  competitive ratio.

We use a smooth approximation of the  $\ell_p$  norm due to [Molinaro \(2017\)](#). Fixing  $\epsilon > 0$  (to be chosen later), for any load vector  $\mathbf{\Lambda} \in \mathbb{R}_{\geq 0}^d$  it is defined by  $\Psi(\mathbf{\Lambda}) := \frac{p}{\epsilon} (\|\mathbf{1} + \frac{\epsilon \mathbf{\Lambda}}{p}\|_p - 1) = \|\frac{p}{\epsilon} \mathbf{1} + \mathbf{\Lambda}\|_p - \frac{p}{\epsilon}$ . In particular, for  $p \rightarrow \infty$ , we have  $\Psi(\mathbf{\Lambda}) \rightarrow \frac{1}{\epsilon} \ln (\sum_j \exp(\epsilon \Lambda_j))$ , which is a common smooth approximation of the  $\ell_\infty$  norm. Sometimes we abbreviate  $\Psi^{(t)} := \Psi(\mathbf{\Lambda}^{(t)})$ .

The following fact follows from the triangle inequality.

**Fact 2.1** (*Additive Approximation, see [Molinaro \(2017\)](#)*) *For any integer  $p \geq 1$  and load  $\mathbf{\Lambda} \in \mathbb{R}_{\geq 0}^d$ , we have  $\|\mathbf{\Lambda}\|_p \leq \Psi(\mathbf{\Lambda}) \leq \|\mathbf{\Lambda}\|_p + \frac{p}{\epsilon} \cdot (\|\mathbf{1}\|_p - 1)$ .*

Note that  $p(\|\mathbf{1}\|_p - 1) \leq \min\{p, \ln d\}$  for all  $p$ . So, this is a meaningful bound even for  $p \rightarrow \infty$ .

The gradient of  $\Psi$  is bounded like that of the function  $\|\cdot\|_p$  (see Fact 2.2). The advantage of  $\Psi$  over  $\|\cdot\|_p$  is that changes in its gradient are also bounded (see Fact 2.3).

**Fact 2.2** (Gradient Norm (Molinari, 2017)) For any integer  $p \geq 1$ ,  $\|\nabla\Psi\|_q \leq 1$  for  $q = (1 - \frac{1}{p})^{-1}$ .

**Fact 2.3** (Gradient Stability (Molinari, 2017)) For any integer  $p \geq 1$ , load  $\Lambda \in \mathbb{R}_{\geq 0}^d$  and load increase  $\mathbf{z} \in [0, 1]^d$ , we have coordinate-wise  $\nabla\Psi(\Lambda + \mathbf{z}) \leq (1 + \epsilon) \cdot \nabla\Psi(\Lambda)$ .

There are several other ways known to get similar smooth approximations, see Nesterov (2005). For our purposes it is convenient that the gradient stability gives a multiplicative approximation.

**Implications for the Surrogate Problem.** Recall that we set the cost of action/expert  $i \in [n]$  in step  $t$  to be  $c_i^{(t)} = \langle C^{(t)} e_i, \nabla\Psi(\Lambda^{(t-1)}) \rangle$ . As an immediate consequence of gradient stability, we can bound the increase of  $\Psi$  by the cost in the surrogate problem as follows.

**Claim 2.4** In any step  $t \in [T]$ , the change in  $\Psi$  is approximated by the cost in the surrogate problem as  $\Psi(\Lambda^{(t)}) - \Psi(\Lambda^{(t-1)}) \leq (1 + \epsilon) \cdot \langle C^{(t)} x^{(t)}, \nabla\Psi(\Lambda^{(t-1)}) \rangle$ .

**Proof** Applying convexity, we get  $\Psi(\Lambda^{(t-1)}) \geq \Psi(\Lambda^{(t)}) + \langle \Lambda^{(t-1)} - \Lambda^{(t)}, \nabla\Psi(\Lambda^{(t)}) \rangle$ . Note that  $\Lambda^{(t-1)} - \Lambda^{(t)} = -C^{(t)} x^{(t)}$ , so  $\Psi(\Lambda^{(t)}) - \Psi(\Lambda^{(t-1)}) \leq \langle C^{(t)} x^{(t)}, \nabla\Psi(\Lambda^{(t)}) \rangle$ . The claim now follows by gradient stability (Fact 2.3) because  $0 \leq (C^{(t)} x^{(t)})_j \leq 1$  for all  $j$ . ■

Applying this bound repeatedly, we can bound the final value of  $\Psi$  on the sequence  $x^{(1)}, \dots, x^{(T)}$  in terms of the cost incurred by the algorithm in the surrogate problem as

$$\Psi(\Lambda^{(T)}) - \Psi(\Lambda^{(0)}) \leq (1 + \epsilon) \cdot \sum_{t=1}^T \langle C^{(t)} x^{(t)}, \nabla\Psi(\Lambda^{(t-1)}) \rangle. \quad (2)$$

As a final remark, we give a bound on the costs of actions in the surrogate problem.

**Claim 2.5** For all  $t \in [T]$  and all  $i \in [n]$ , we have  $0 \leq c_i^{(t)} \leq \|\mathbf{1}\|_p$ .

**Proof** By Hölder's inequality, we have for all  $t \in [T]$  and all  $i \in [n]$  that  $c_i^{(t)} = \langle C^{(t)} e_i, \nabla\Psi(\Lambda^{(t-1)}) \rangle$  is at most  $\|C^{(t)} e_i\|_p \cdot \|\nabla\Psi(\Lambda^{(t-1)})\|_q$ , where  $q = (1 - \frac{1}{p})^{-1}$ . From Fact 2.2,  $\|\nabla\Psi(\Lambda^{(t-1)})\|_q \leq 1$ . Furthermore,  $0 \leq C_{i,j}^{(t)} \leq 1$  for all  $i$  and  $j$ , resulting in  $\|C^{(t)} e_i\|_p \leq \|\mathbf{1}\|_p$ . ■

### 3. Load Balancing

Recall that in the load balancing problem there is a sequence of cost matrices  $C^{(1)}, \dots, C^{(T)} \in [0, 1]^{d \times n}$ , which are either drawn i.i.d. from an unknown distribution or defined by an adversary. Our algorithm chooses  $x^{(1)}, \dots, x^{(T)}$  to minimize  $\|\Lambda^{(T)}\|_p$ , where  $\Lambda^{(t)} = \sum_{s=1}^t C^{(s)} \cdot x^{(s)}$ . The point of comparison is any fixed action benchmark  $x^* \in \Delta_n$ , and  $\text{OPT}_{\text{OLVC}}$  is defined in Eq. (1).

To solve the load balancing problem, we can apply any no-regret algorithm that is able to cope with an adaptive adversary in the surrogate problem. This ensures that in the surrogate game in hindsight it would not have been much better to always have chosen a different action. Due to Claim 2.5, we have  $0 \leq \frac{1}{\|\mathbf{1}\|_p} c_i^{(t)} \leq 1$ . As a consequence,  $\sum_{t=1}^T c^{(t)} x^{(t)} \leq \min_{x \in \Delta_n} \sum_{t=1}^T c^{(t)} x + \|\mathbf{1}\|_p \cdot$

REGRET, where REGRET denotes the regret bound for the respective one-dimensional problem in which the costs are bounded between 0 and 1. So, for any no-regret algorithm,  $\text{REGRET} = o(T)$ .

Note that depending on the feedback, different online learning algorithms can be applied. If we get to know  $C^{(t)}$  entirely, we can compute the entire vector  $c^{(t)}$  and pass it on to the algorithm, which correspond to experts feedback. In this case, Hedge (Freund and Schapire, 1999) guarantees  $\text{REGRET} = O(\sqrt{T \log n})$ . If we only get to know  $C^{(t)}x^{(t)}$ , we can still compute  $c^{(t)}x^{(t)}$ , which corresponds to bandits feedback. Now, Exp3.P (Auer et al., 2002) guarantees  $\text{REGRET} = \tilde{O}(\sqrt{T \cdot n})$ .

Combining this regret guarantee with Claim 2.4, we can bound the final value of  $\Psi$  in terms of the cost of any benchmark solution  $x^*$  in the surrogate problem. In particular, letting  $x$  denote the choice of actions  $x^*$  that minimizes  $\|\sum_{t=1}^T C^{(t)}x^*\|_p$ , we have

$$\begin{aligned} \Psi(\mathbf{\Lambda}^{(T)}) &\leq \Psi(\mathbf{\Lambda}^{(0)}) + (1 + \epsilon) \cdot \sum_{t=1}^T \langle C^{(t)}x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \\ &\leq (1 + \epsilon) \cdot \sum_{t=1}^T \langle C^{(t)}x^*, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle + \|\mathbf{1}\|_p \cdot \text{REGRET} + \Psi(\mathbf{\Lambda}^{(0)}). \end{aligned} \quad (3)$$

The difficulty is to relate the cost  $\sum_{t=1}^T \langle C^{(t)}x^*, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle$  incurred by  $x^*$  in the surrogate game to the  $\ell_p$  norm of the load it would have generated. Importantly, we cannot apply Claim 2.4 for such a comparison because the costs in the surrogate game are still defined via loads  $\mathbf{\Lambda}^{(0)}, \dots, \mathbf{\Lambda}^{(T-1)}$  generated by the algorithm and not by  $x^*$ .

In the stochastic setting, we can apply independence between steps. This lets us derive that  $\mathbb{E}[\|\mathbf{\Lambda}^{(T)}\|_p] \leq (1 + \epsilon) \cdot \text{OPT}_{\text{OLVC}} + \|\mathbf{1}\|_p \cdot \text{REGRET} + 2\frac{p}{\epsilon} \cdot (\|\mathbf{1}\|_p - 1)$  in §3.1. In the adversarial setting we have to make much more explicit use of the definition of  $\Psi$ , which lets us derive  $\|\mathbf{\Lambda}^{(T)}\|_p = O(\min\{p, \log d\} \cdot \text{OPT}_{\text{OLVC}} + \|\mathbf{1}\|_p \cdot \text{REGRET})$  in §3.2.

### 3.1. Stochastic Arrivals

We first consider the stochastic setting where in each time step the cost matrix  $C^{(t)} \in [0, 1]^{d \times n}$  is sampled i.i.d. from some unknown distribution. For any benchmark  $x^*$ , we define  $\text{OPT}_{\text{OLVC}} := \mathbb{E}[\|\sum_{t=1}^T C^{(t)} \cdot x^*\|_p]$ . We derive the following bound on the  $\ell_p$  norm of the load.

**Theorem 1** *For stochastic load balancing, any no-regret algorithm applied to the surrogate problem guarantees a load bound of  $\mathbb{E}[\|\mathbf{\Lambda}^{(T)}\|_p] \leq (1 + \epsilon) \cdot \text{OPT}_{\text{OLVC}} + \|\mathbf{1}\|_p \cdot \text{REGRET} + 2\frac{p}{\epsilon} \cdot (\|\mathbf{1}\|_p - 1)$ , where REGRET is the regret of one-dimensional online learning.*

For the special case of  $\ell_\infty$  norm, we get a bound of  $\mathbb{E}[\|\mathbf{\Lambda}^{(T)}\|_\infty] \leq (1 + \epsilon) \cdot \text{OPT}_{\text{OLVC}} + \text{REGRET} + 2\frac{\ln d}{\epsilon}$  because  $p(\|\mathbf{1}\|_p - 1) \rightarrow \ln d$  as  $p \rightarrow \infty$ . At the heart of Theorem 1 is the following claim (which is inspired from Molinaro (2017)) relating the surrogate cost of  $x^*$  to  $\text{OPT}_{\text{OLVC}}$ .

**Claim 3.1** *If  $C^{(1)}, \dots, C^{(T)}$  are independent, identically distributed random variables,*

$$\mathbb{E}[\langle C^{(t)}x^*, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle] \leq \left\| \mathbb{E}_{C^{(t)}}[C^{(t)}x^*] \right\|_p \leq \frac{\text{OPT}_{\text{OLVC}}}{T}.$$

**Proof** Since  $C^{(1)}, \dots, C^{(T)}$  are independent, we can first take the expectation only over  $C^{(t)}$ . As  $\nabla \Psi(\mathbf{\Lambda}^{(t-1)})$  only depends on the past actions, this yields  $\mathbb{E}_{C^{(t)}}[\langle C^{(t)}x^*, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle] = \langle \mathbb{E}_{C^{(t)}}[C^{(t)}x^*], \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle$ . Again using independence of  $C^{(1)}, \dots, C^{(T)}$ , we get  $\mathbb{E}_{C^{(t)}}[C^{(t)}x^*] = \mathbb{E}[C^{(t)}x^*]$ . Applying Hölder's inequality and using  $\|\nabla \Psi(\mathbf{\Lambda}^{(t-1)})\|_q \leq 1$ , we get for any  $\mathbf{\Lambda}^{(t-1)}$  that

$$\langle \mathbb{E}[C^{(t)}x^*], \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \leq \|\mathbb{E}[C^{(t)}x^*]\|_p \cdot \|\nabla \Psi(\mathbf{\Lambda}^{(t-1)})\|_q \leq \|\mathbb{E}[C^{(t)}x^*]\|_p.$$

So, in combination  $\mathbb{E}_{C^{(t)}}[\langle C^{(t)}x^*, \nabla\Psi(\Lambda^{(t-1)}) \rangle] \leq \|\mathbb{E}[C^{(t)}x^*]\|_p$ . The right-hand side is not a random variable. So taking expectation over  $C^{(1)}, \dots, C^{(t-1)}$  proves the first inequality.

For the second inequality, since  $\|\cdot\|_p$  is a convex function it satisfies  $\mathbb{E}[\|\cdot\|_p] \geq \|\mathbb{E}[\cdot]\|_p$ . Hence,

$$\text{OPT}_{\text{OLVC}} = \mathbb{E}\left[\left\|\sum_{t=1}^T C^{(t)} \cdot x^*\right\|_p\right] \geq \left\|\mathbb{E}\left[\sum_{t=1}^T C^{(t)} \cdot x^*\right]\right\|_p = T \cdot \left\|\mathbb{E}[C^{(t)} \cdot x^*]\right\|_p,$$

which completes the proof of Claim 3.1.  $\blacksquare$

Now the proof of Theorem 1 is straightforward.

**Proof** (of Theorem 1) Combining Claim 3.1 with Eq. (3) and taking expectations,

$$\begin{aligned} \mathbb{E}[\Psi^{(T)}] &\leq (1 + \epsilon) \cdot \sum_{t \leq T} \mathbb{E}[\langle C^{(t)}x^*, \nabla\Psi(\Lambda^{(t-1)}) \rangle] + \|\mathbf{1}\|_p \cdot \text{REGRET} + \Psi^{(0)} \\ &\leq (1 + \epsilon) \cdot T \cdot \frac{\text{OPT}_{\text{OLVC}}}{T} + \|\mathbf{1}\|_p \cdot \text{REGRET} + \Psi^{(0)} \\ &= (1 + \epsilon) \cdot \text{OPT}_{\text{OLVC}} + \|\mathbf{1}\|_p \cdot \text{REGRET} + \frac{p}{\epsilon} \cdot (\|\mathbf{1}\|_p - 1). \end{aligned}$$

The theorem now follows by the additive approximation property of  $\Psi$  (Fact 2.1).  $\blacksquare$

### 3.2. Adversarial Arrivals

Now we consider the adversarial setting where the cost matrices  $C^{(1)}, \dots, C^{(T)} \in [0, 1]^{d \times n}$  are defined by an adversary. Analogous to the stochastic setting, for any benchmark  $x^*$ , we define  $\text{OPT}_{\text{OLVC}} := \left\|\sum_{t=1}^T C^{(t)} \cdot x^*\right\|_p$ . The following is our main result for adversarial load balancing.

**Theorem 2** *For adversarial load balancing  $\text{OLVC}_p$ , any no-regret algorithm applied to the surrogate problem guarantees a load bound of  $\|\Lambda^{(T)}\|_p \leq O(\min\{p, \log d\} \cdot \text{OPT}_{\text{OLVC}}) + O(\|\mathbf{1}\|_p \cdot \text{REGRET})$ , where REGRET is the regret of the one-dimensional online learning problem.*

Like in the stochastic setting, we have to bound the loss of  $x^*$  in the surrogate game. Claim 3.1 uses the stochastic assumptions very heavily and does not hold in this more general case. As a replacement, we show the following lemma. It disentangles the loss of  $x^*$  in the surrogate game into a factor depending on the actual load that  $x^*$  generates and the cost that the algorithm incurs.

**Lemma 3** *For any action vector  $x^* \in \Delta_n$  and  $\text{OPT}_{\text{OLVC}} := \left\|\sum_{t=1}^T C^{(t)}x^*\right\|_p$ , we have*

$$\sum_{t=1}^T \langle C^{(t)}x^*, \nabla\Psi(\Lambda^{(t-1)}) \rangle \leq \left(\exp\left(\frac{\epsilon \text{OPT}_{\text{OLVC}}}{\|\mathbf{1}\|_p}\right) - 1\right) \left( (1 + \epsilon) \sum_{t=1}^T \langle C^{(t)}x^{(t)}, \nabla\Psi(\Lambda^{(t-1)}) \rangle + \frac{\|\mathbf{1}\|_p}{\epsilon} \right).$$

Given Lemma 3, the remaining proof of Theorem 2 is rather straightforward.

**Proof** (of Theorem 2) The one-dimensional regret bound in combination with Lemma 3 gives

$$\begin{aligned} \sum_{t=1}^T \langle C^{(t)}x^{(t)}, \nabla\Psi(\Lambda^{(t-1)}) \rangle &\leq \sum_{t=1}^T \langle C^{(t)}x^*, \nabla\Psi(\Lambda^{(t-1)}) \rangle + \|\mathbf{1}\|_p \cdot \text{REGRET} \\ &\leq \left(\exp\left(\frac{\epsilon \text{OPT}_{\text{OLVC}}}{\|\mathbf{1}\|_p}\right) - 1\right) \cdot \left( (1 + \epsilon) \left( \sum_{t=1}^T \langle C^{(t)}x^*, \nabla\Psi(\Lambda^{(t-1)}) \rangle \right) + \frac{\|\mathbf{1}\|_p}{\epsilon} \right) + \|\mathbf{1}\|_p \cdot \text{REGRET}. \end{aligned}$$



Setting  $\epsilon = \min \left\{ 1, \frac{\|\mathbf{1}\|_p}{5 \cdot \text{OPT}_{\text{OLVC}}} \right\}$  gives  $\left( \exp \left( \frac{\epsilon \text{OPT}_{\text{OLVC}}}{\|\mathbf{1}\|_p} \right) - 1 \right) \leq \exp\left(\frac{1}{5}\right) - 1 \leq \frac{1}{4}$ . Rearranging,

$$\left( 1 - \frac{1+\epsilon}{4} \right) \sum_{t=1}^T \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \leq \frac{\|\mathbf{1}\|_p}{4\epsilon} + \|\mathbf{1}\|_p \cdot \text{REGRET},$$

which implies  $\sum_{t=1}^T \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \leq \frac{\|\mathbf{1}\|_p}{2\epsilon} + 2\|\mathbf{1}\|_p \cdot \text{REGRET}$ .

Finally, by Claim 2.4, we have  $\Psi^{(T)} \leq \Psi^{(0)} + (1+\epsilon) \sum_{t=1}^T \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle$ . Now using Fact 2.1 and  $\Psi^{(0)} = \frac{p}{\epsilon} (\|\mathbf{1}\|_p - 1)$  gives

$$\begin{aligned} \|\mathbf{\Lambda}^{(T)}\|_p &\leq \Psi^{(T)} \leq \frac{p}{\epsilon} (\|\mathbf{1}\|_p - 1) + (1+\epsilon) \frac{\|\mathbf{1}\|_p}{2\epsilon} + 2(1+\epsilon)\|\mathbf{1}\|_p \cdot \text{REGRET} \\ &\leq 5 \left( 1 + p \frac{\|\mathbf{1}\|_p - 1}{\|\mathbf{1}\|_p} \right) \text{OPT}_{\text{OLVC}} + 4\|\mathbf{1}\|_p \text{REGRET} + p(\|\mathbf{1}\|_p - 1) + \|\mathbf{1}\|_p \\ &= O(\min\{p, \log d\} \cdot \text{OPT}_{\text{OLVC}}) + O(\|\mathbf{1}\|_p \cdot \text{REGRET}), \end{aligned}$$

which proves Theorem 2. ■

In the rest of the section we prove the missing Lemma 3. We first rewrite  $\Psi$  in a different form, namely for all  $\mathbf{\Lambda} \in \mathbb{R}_{\geq 0}^d$ ,

$$\Psi(\mathbf{\Lambda}) = \frac{p}{\epsilon} \left( (\Phi(\mathbf{\Lambda}))^{1/p} - 1 \right) \quad \text{where} \quad \Phi(\mathbf{\Lambda}) := \sum_{j=1}^d \left( 1 + \frac{\epsilon}{p} \Lambda_j \right)^p.$$

For intuition, note that  $p \rightarrow \infty$  implies  $\Phi(\mathbf{\Lambda}) \rightarrow \sum_j \exp(\epsilon \Lambda_j)$ , a commonly used potential. So,

$$\left( \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \right)_j = \left( 1 + \frac{\epsilon}{p} \Lambda_j^{(t-1)} \right)^{p-1} \cdot (\Phi(\mathbf{\Lambda}^{(t-1)}))^{-1/q}. \quad (4)$$

We will need the following property of the function  $\Phi$ , which follows from the monotonicity of  $\nabla \Phi$  and is a slight variant of Lemma 3.1 in Caragiannis (2008). We defer its proof to §D.

**Lemma 4** (Caragiannis (2008)) *For any integer  $k \geq 1$ , any  $\mathbf{\Lambda} \in \mathbb{R}_{\geq 0}^d$ , and any  $z^{(i)} \in \mathbb{R}_{\geq 0}^d$  for  $i \in [k]$ , we have  $\sum_{i=1}^k \left( \Phi(\mathbf{\Lambda} + z^{(i)}) - \Phi(\mathbf{\Lambda}) \right) \leq \Phi\left(\mathbf{\Lambda} + \sum_{i=1}^k z^{(i)}\right) - \Phi(\mathbf{\Lambda})$ .*

We now have the ingredients to prove Lemma 3.

**Proof** (of Lemma 3) Using Eq. (4) and convexity of  $\Phi$ , we have

$$\begin{aligned} \sum_{t=1}^T \langle C^{(t)} x^*, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle &= \sum_{t=1}^T \left\langle C^{(t)} x^*, \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(t-1)})} \nabla \Phi(\mathbf{\Lambda}^{(t-1)}) \right\rangle \\ &\leq \sum_{t=1}^T \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(t-1)})} \left( \Phi(\mathbf{\Lambda}^{(t-1)} + C^{(t)} x^*) - \Phi(\mathbf{\Lambda}^{(t-1)}) \right). \end{aligned}$$

For  $i \in [T]$ , define

$$a_i := \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(i-1)})} - \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(i)})} \quad \text{and} \quad a_T := \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(T)})}. \quad (5)$$

This implies  $\frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(t-1)})} = \sum_{i=t}^{T-1} \left( \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(i-1)})} - \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(i)})} \right) + \frac{1}{\epsilon \Phi^{1/q}(\mathbf{\Lambda}^{(T)})} = \sum_{i=t}^T a_i$ . So, the above inequality can be rewritten as

$$\begin{aligned} \sum_{t=1}^T \langle C^{(t)} x^*, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle &\leq \sum_{t=1}^T \sum_{i=t}^T a_i \cdot \left( \Phi(\mathbf{\Lambda}^{(t-1)} + C^{(t)} x^*) - \Phi(\mathbf{\Lambda}^{(t-1)}) \right) \\ &= \sum_{i=1}^T a_i \cdot \sum_{t=1}^i \left( \Phi(\mathbf{\Lambda}^{(t-1)} + C^{(t)} x^*) - \Phi(\mathbf{\Lambda}^{(t-1)}) \right) \\ &\leq \sum_{i=1}^T a_i \cdot \sum_{t=1}^i \left( \Phi(\mathbf{\Lambda}^{(i-1)} + C^{(t)} x^*) - \Phi(\mathbf{\Lambda}^{(i-1)}) \right), \end{aligned}$$

where the last inequality uses monotonicity of  $\nabla\Phi$  in every component. Now applying Lemma 4,

$$\sum_{t=1}^T \langle C^{(t)}x^*, \nabla\Psi(\Lambda^{(t-1)}) \rangle \leq \sum_{i=1}^T a_i \cdot \left( \Phi\left(\sum_{t=1}^i C^{(t)}x^* + \Lambda^{(i-1)}\right) - \Phi(\Lambda^{(i-1)}) \right). \quad (6)$$

To simplify the right-hand side, we use the definition of  $\Phi$  and Minkowski's inequality to get

$$\begin{aligned} \Phi^{1/p}\left(\sum_{t=1}^i C^{(t)}x^* + \Lambda^{(i-1)}\right) &= \left\| \mathbf{1} + \frac{\epsilon}{p} \left(\sum_{t=1}^i C^{(t)}x^* + \Lambda^{(i-1)}\right) \right\|_p \\ &\leq \left\| \mathbf{1} + \frac{\epsilon}{p} \Lambda^{(i-1)} \right\|_p + \left\| \frac{\epsilon}{p} \cdot \sum_{t=1}^i C^{(t)}x^* \right\|_p. \end{aligned}$$

To further simplify, note that

$$\left\| \frac{\epsilon}{p} \cdot \sum_{t=1}^i C^{(t)}x^* \right\|_p \leq \frac{\epsilon}{p} \text{OPT}_{\text{OLVC}} \leq \frac{\epsilon}{p} \text{OPT}_{\text{OLVC}} \cdot \frac{\Phi^{1/p}(\Lambda^{(i-1)})}{d^{1/p}},$$

where the last step uses that  $\Phi(\Lambda^{(i-1)}) \geq d$ . Since  $\left\| \mathbf{1} + \frac{\epsilon}{p} \Lambda^{(i-1)} \right\|_p = \Phi^{1/p}(\Lambda^{(i-1)})$ , this overall simplifies the above expression to

$$\begin{aligned} \Phi\left(\sum_{t=1}^i C^{(t)}x^* + \Lambda^{(i-1)}\right) &\leq \left( \Phi(\Lambda^{(i-1)})^{1/p} + \frac{\epsilon \text{OPT}_{\text{OLVC}}}{pd^{1/p}} \Phi(\Lambda^{(i-1)})^{1/p} \right)^p \\ &= \Phi(\Lambda^{(i-1)}) \left( 1 + \frac{\epsilon \text{OPT}_{\text{OLVC}}}{pd^{1/p}} \right)^p \leq \Phi(\Lambda^{(i-1)}) \cdot \exp\left(\frac{\epsilon \text{OPT}_{\text{OLVC}}}{d^{1/p}}\right). \end{aligned}$$

So, in combination with Eq. (6), we get

$$\sum_{t=1}^T \langle C^{(t)}x^*, \nabla\Psi(\Lambda^{(t-1)}) \rangle \leq \left( \exp\left(\frac{\epsilon \text{OPT}_{\text{OLVC}}}{d^{1/p}}\right) - 1 \right) \cdot \sum_{i=1}^T a_i \cdot \Phi(\Lambda^{(i-1)}),$$

which finishes the proof of Lemma 3 by applying the following Claim 3.2.  $\blacksquare$

**Claim 3.2** For  $a_i$ 's defined in Eq. (5), we have

$$\sum_{i=1}^T a_i \cdot \Phi(\Lambda^{(i-1)}) \leq (1 + \epsilon) \left( \sum_{t=1}^T \langle C^{(t)}x^{(t)}, \nabla\Psi(\Lambda^{(t-1)}) \rangle \right) + \frac{d^{1/p}}{\epsilon}.$$

**Proof** We start by expanding the left-hand side to get

$$\begin{aligned} \sum_{i=1}^T a_i \cdot \Phi(\Lambda^{(i-1)}) &= \sum_{i=1}^T \left( \frac{1}{\epsilon \Phi^{1/q}(\Lambda^{(i-1)})} - \frac{1}{\epsilon \Phi^{1/q}(\Lambda^{(i)})} \right) \cdot \Phi(\Lambda^{(i-1)}) \\ &= \sum_{i=2}^T \left( \frac{\Phi(\Lambda^{(i-1)}) - \Phi(\Lambda^{(i-2)})}{\epsilon \Phi^{1/q}(\Lambda^{(i-1)})} \right) + \frac{\Phi(\Lambda^{(0)})}{\epsilon \Phi^{1/q}(\Lambda^{(0)})} - \frac{\Phi(\Lambda^{(T-1)})}{\epsilon \Phi^{1/q}(\Lambda^{(T)})} \\ &\leq \sum_{i=2}^T \left( \frac{\Phi(\Lambda^{(i-2)} + C^{(i-1)}x^{(i-1)}) - \Phi(\Lambda^{(i-2)})}{\epsilon \Phi^{1/q}(\Lambda^{(i-1)})} \right) + \frac{d^{1/p}}{\epsilon}. \end{aligned}$$

Since  $\Phi(\cdot)$  is a convex function,  $\Phi(\Lambda^{(i-2)} + C^{(i-1)}x^{(i-1)}) - \Phi(\Lambda^{(i-2)}) \leq \langle \nabla\Phi(\Lambda^{(i-1)}), C^{(i-1)}x^{(i-1)} \rangle$ .

Moreover, recall from Eq. (4) that  $\nabla\Psi^{(i-1)} = \frac{\nabla\Phi(\Lambda^{(i-1)})}{\epsilon \Phi^{1/q}(\Lambda^{(i-1)})}$ . Thus, we get

$$\begin{aligned} \sum_{i=1}^T a_i \cdot \Phi(\Lambda^{(i-1)}) &\leq \sum_{i=2}^T \langle \nabla\Psi^{(i-1)}, C^{(i-1)}x^{(i-1)} \rangle + \frac{d^{1/p}}{\epsilon} \\ &\leq (1 + \epsilon) \sum_{i=2}^T \langle \nabla\Psi^{(i-2)}, C^{(i-1)}x^{(i-1)} \rangle + \frac{d^{1/p}}{\epsilon} \\ &= (1 + \epsilon) \left( \sum_{t=1}^{T-1} \langle C^{(t)}x^{(t)}, \nabla\Psi(\Lambda^{(t-1)}) \rangle \right) + \frac{d^{1/p}}{\epsilon}, \end{aligned}$$

where the last inequality uses the stability of gradient of  $\Psi$  (Fact 2.3).  $\blacksquare$

#### 4. Adversarial Bandits with Knapsacks

In this section we consider  $\text{BWK}_p$  as defined in §1.2. The goal is to maximize the total collected reward while ensuring that the  $\ell_p$  norm of the total cost vector stays within the budget  $B$ . Recall, for  $\ell_\infty$  norm this captures in the stochastic case the setting of [Badanidiyuru et al. \(2018\)](#) and in the adversarial case the setting of [Immorlica et al. \(2019\)](#).

The high-level plan of our proofs is to use the ideas from §3 to reduce the  $d$ -dimensional budget constraint to a single-dimensional budget constraint, and then to Lagrangify this single budget constraint. In this section we present how to execute this plan in the adversarial case, where the proof is easier as we can afford losing a constant-multiplicative factor in the approximation. In §A we discuss the stochastic case where we need more work and use a different norm.

To deal with the multi-dimensional budget constraint, we follow the approach described in §2 and define  $\Psi(\mathbf{\Lambda}) = \frac{p}{\epsilon} (\|\mathbf{1} + \frac{\epsilon \mathbf{\Lambda}}{p}\|_p - 1)$ . We again consider a surrogate one-dimensional online learning game with  $n$  actions. This time, the algorithm's goal is to maximize the following suitably defined reward function. In time step  $t$ , the *Lagrangian reward* of the  $i$ -th action/expert is set to

$$\mathcal{R}_i^{(t)} = \begin{cases} r^{(t)} e_i - \lambda \cdot \langle C^{(t)} e_i, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle & \text{if } \|\mathbf{\Lambda}^{(t-1)}\|_p \leq B \\ 0 & \text{otherwise,} \end{cases}$$

where  $r^{(t)}$  is the reward vector from  $\text{BWK}_p$  problem,  $\mathbf{\Lambda}^{(t)} = \sum_{s=1}^t C^{(s)} \cdot x^{(s)}$  is the load after time  $t$ , and  $\nabla \Psi(\mathbf{\Lambda}^{(t-1)})$  is the gradient of  $\Psi$ . We will use

$$\lambda := \frac{\sum_{t=1}^{\tau^*} r^{(t)} \cdot x^*}{2B} = \frac{\text{OPT}_{\text{BWK}}}{2B} \quad \text{and} \quad \epsilon := \frac{2p(\|\mathbf{1}\|_p - 1)}{B},$$

where  $\tau^*$  is the time at which  $x^*$  runs out of budget, otherwise  $\tau^* = T$ . Note that the rewards are always zero after the total cost/load vector exceeds budget  $B$ . Moreover, the algorithm can implement this game online since we are assuming the adversary is adaptive (non-oblivious).

**Claim 4.1** *The reward of each expert per time step is bounded in magnitude by  $(\lambda \cdot \|\mathbf{1}\|_p + 1)$ .*

**Proof** By Hölder's inequality, we have  $\langle C^{(t)} e_i, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \leq \|C^{(t)} e_i\|_p \cdot \|\nabla \Psi(\mathbf{\Lambda}^{(t-1)})\|_q \leq \|\mathbf{1}\|_p$ , where the last step uses  $\|\nabla \Psi(\mathbf{\Lambda}^{(t-1)})\|_q \leq 1$  by Fact 2.2 and  $\|C^{(t)} e_i\|_p \leq \|\mathbf{1}\|_p$ .

Consequently  $|\mathcal{R}_i^{(t)}| \leq |r^{(t)} e_i| + |\lambda \langle C^{(t)} e_i, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle| \leq \lambda \cdot \|\mathbf{1}\|_p + 1$ .  $\blacksquare$

Let  $\text{REGRET}_t$  denote the regret of a 1-dimensional online learning algorithm against an adaptive adversary at time  $t \in \{1, \dots, T\}$ .

Recall,  $x^*$  is defined to be a benchmark solution satisfying  $\|\sum_{t=1}^T C^{(t)} x^*\|_p \leq B$ . Our point of comparison will be a scaled-down version  $x' := \frac{1}{\alpha} x^*$ , where  $\alpha := 5p \frac{\|\mathbf{1}\|_p - 1}{\|\mathbf{1}\|_p}$  (note  $\alpha \rightarrow 5 \ln d$  as  $p \rightarrow \infty$ ). The intuition for scaling  $x^*$  is that for adversarial OLVC we only know from Theorem 2 that the budget is not exceeded multiplicatively by more than a factor of  $p$ . So to *strictly* satisfy the budget constraints in  $\text{BWK}_p$ , we need to scale-down the solution. The following is the main result.

**Theorem 5** *If  $B \geq 2p(\|\mathbf{1}\|_p - 1)$  then for adversarial  $\text{BWK}_p$  the algorithm gets a reward*

$$\sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} \geq \frac{1}{20 \min\{p, \ln d\}} \text{OPT}_{\text{BWK}} - O\left(\frac{\text{OPT}_{\text{BWK}} \cdot \|\mathbf{1}\|_p}{B}\right) \cdot \text{REGRET}.$$

**Proof** We distinguish the analysis in two cases: Either the algorithm stops before time  $\tau^*$  (recall, this is time at which  $x^*$  runs out of budget, otherwise  $\tau^* = T$ ) or it stays within budget until  $\tau^*$ .

**Case 1: Algorithm stops before  $\tau^*$**   $\|\mathbf{\Lambda}^{(\tau^*)}\|_p > B$ .

In this case, there is some time  $\tau \leq \tau^*$  at which  $\|\mathbf{\Lambda}^{(\tau)}\|_p > B$  for the first time. This means that  $\tau$  is the last round before the algorithm stops and  $\mathcal{R}^{(t)} = 0$  for all  $t > \tau$ . Since Claim 4.1 bounds the loss of each action/expert, we get  $O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*}$  bounds the Lagrangian reward of the algorithm compared to always playing the null action, which has Lagrangian reward 0. This gives

$$\sum_{t=1}^{\tau} \mathcal{R}^{(t)} \cdot x^{(t)} = \sum_{t=1}^{\tau^*} \mathcal{R}^{(t)} \cdot x^{(t)} \geq 0 - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*}.$$

Using the definition of  $\mathcal{R}^{(t)}$ , this implies

$$\sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} \geq \lambda \sum_{t=1}^{\tau} \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*}$$

We have  $\Psi(\mathbf{\Lambda}^{(\tau)}) - \Psi(0) \geq \|\mathbf{\Lambda}^{(\tau)}\|_p - \frac{p}{\epsilon} \cdot (\|\mathbf{1}\|_p - 1) = \frac{B}{2}$ . So, by Claim 2.4,

$$\sum_{t=1}^{\tau} \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \geq \frac{1}{1+\epsilon} (\Psi(\mathbf{\Lambda}^{(\tau)}) - \Psi(0)) \geq \frac{B}{2(1+\epsilon)}.$$

This implies

$$\sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} \geq \frac{\lambda \cdot B}{2(1+\epsilon)} - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*} = \frac{\text{OPT}_{\text{BwK}}}{4(1+\epsilon)} - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*}.$$

**Case 2: Algorithm stays within budget till  $\tau^*$**   $\|\mathbf{\Lambda}^{(\tau^*)}\|_p \leq B$ .

In this case,  $\|\mathbf{\Lambda}^{(t)}\|_p \leq B$  for all relevant  $t$ . This time, we use that  $\text{REGRET}_{\tau^*}$  bounds the Lagrangian reward compared to always playing  $x'$ . That is,  $\sum_{t=1}^{\tau^*} \mathcal{R}^{(t)} \cdot x^{(t)} \geq \sum_{t=1}^{\tau^*} \mathcal{R}^{(t)} \cdot x' - \text{REGRET}_{\tau^*}$ , and therefore by the definition of  $\mathcal{R}^{(t)}$ ,

$$\begin{aligned} \sum_{t=1}^{\tau^*} r^{(t)} \cdot x^{(t)} &\geq \sum_{t=1}^{\tau^*} (r^{(t)} \cdot x' - \lambda \cdot \langle C^{(t)} x', \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle) + \lambda \cdot \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle \\ &\quad - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*}. \end{aligned} \quad (7)$$

Now apply Lemma 3 to  $x'$ , noting that  $\|\sum_{t=1}^{\tau^*} C^{(t)} x'\|_p \leq \frac{B}{\alpha}$ . So,

$$\begin{aligned} \sum_{t=1}^{\tau^*} \langle C^{(t)} x', \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle &\leq \left( \exp\left(\frac{\epsilon}{\|\mathbf{1}\|_p} \frac{B}{\alpha}\right) - 1 \right) \left( (1+\epsilon) \sum_{t=1}^{\tau^*} \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle + \frac{\|\mathbf{1}\|_p}{\epsilon} \right) \\ &= (e^{2/5} - 1) \left( (1+\epsilon) \sum_{t=1}^{\tau^*} \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle + \frac{\|\mathbf{1}\|_p}{\epsilon} \right) \\ &\leq \sum_{t=1}^{\tau^*} \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle + \frac{\|\mathbf{1}\|_p}{2\epsilon}, \end{aligned}$$

where the last step uses that  $\epsilon \leq 1$ . We can rewrite this inequality as

$$\begin{aligned} -\lambda \sum_{t=1}^{\tau^*} \langle C^{(t)} x', \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle + \lambda \sum_{t=1}^{\tau^*} \langle C^{(t)} x^{(t)}, \nabla \Psi(\mathbf{\Lambda}^{(t-1)}) \rangle &\geq -\lambda \frac{\|\mathbf{1}\|_p}{2\epsilon} \\ &= -\lambda \frac{B}{4p} \frac{\|\mathbf{1}\|_p}{\|\mathbf{1}\|_p - 1} \geq -\frac{1}{8p} \frac{\|\mathbf{1}\|_p}{\|\mathbf{1}\|_p - 1} \sum_{t=1}^{\tau^*} r^{(t)} \cdot x^* \end{aligned}$$

by our choice of  $\lambda = \frac{\sum_{t=1}^{\tau^*} r^{(t)} \cdot x^*}{2B}$  and  $\epsilon = \frac{2p(\|\mathbf{1}\|_p - 1)}{B}$ . Substituting this in Eq. (7) gives

$$\begin{aligned} \sum_{t=1}^{\tau^*} r^{(t)} \cdot x^{(t)} &\geq \sum_{t=1}^{\tau^*} \frac{1}{\alpha} \langle r^{(t)}, x^* \rangle - \frac{1}{8p} \frac{\|\mathbf{1}\|_p}{\|\mathbf{1}\|_p - 1} \sum_{t=1}^{\tau^*} r^{(t)} \cdot x^* - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*} \\ &\geq \frac{1}{20 \min\{p, \ln d\}} \sum_{t=1}^{\tau^*} r^{(t)} \cdot x^* - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}_{\tau^*} \end{aligned}$$

because  $\alpha = 5p \frac{\|\mathbf{1}\|_p - 1}{\|\mathbf{1}\|_p}$ , which completes the proof.  $\blacksquare$

## Acknowledgments

We are thankful to the authors of [Immorlica et al. \(2019\)](#) for explaining us their results. We thank the anonymous reviewers of COLT 2020 for helping us improve the presentation of the paper.

## References

- Jacob D. Abernethy, Peter L. Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of COLT*, pages 27–46, 2011.
- Shipra Agrawal and Nikhil R. Devanur. Bandits with concave rewards and convex knapsacks. In *ACM Conference on Economics and Computation*, pages 989–1006, 2014.
- Shipra Agrawal and Nikhil R. Devanur. Linear contextual bandits with knapsacks. In *Advances in Neural Information Processing Systems*, pages 3450–3458, 2016.
- Shipra Agrawal, Nikhil R. Devanur, and Lihong Li. An efficient algorithm for contextual bandits with knapsacks, and an extension to concave objectives. In *Proceedings of COLT*, 2016.
- Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multi-armed bandit problem. *SIAM J. Comput.*, 32(1):48–77, 2002.
- Baruch Awerbuch, Yossi Azar, Edward F. Grove, Ming-Yang Kao, P. Krishnan, and Jeffrey Scott Vitter. Load balancing in the  $l_p$  norm. In *Proceedings of FOCS*, pages 383–391, 1995.
- Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. Bandits with knapsacks. *J. ACM*, 65(3):13:1–13:55, 2018.
- David Blackwell et al. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
- Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 2005.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- Niv Buchbinder, Joseph Seffi Naor, et al. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends® in Theoretical Computer Science*, 3(2–3), 2009.
- Ioannis Caragiannis. Better bounds for online load balancing on unrelated machines. In *Proceedings of SODA*, pages 972–981, 2008.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.

- Sungjin Im, Nathaniel Kell, Janardhan Kulkarni, and Debmalya Panigrahi. Tight bounds for online vector scheduling. *SIAM J. Comput.*, 48(1):93–121, 2019.
- Nicole Immorlica, Karthik Abinav Sankararaman, Robert E. Schapire, and Aleksandrs Slivkins. Adversarial bandits with knapsacks. In *Proceedings of FOCS*, pages 202–219, 2019.
- Sham M Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106, 2009.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. In *30th Annual Symposium on Foundations of Computer Science*, 1989.
- Aranyak Mehta. Online matching and ad allocation. *Theoretical Computer Science*, 8(4), 2012.
- Marco Molinaro. Online and random-order load balancing simultaneously. In *Proceedings of SODA*, pages 1638–1650, 2017.
- Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
- Anshuka Rangi, Massimo Franceschetti, and Long Tran-Thanh. Unifying the stochastic and the adversarial bandits with knapsack. In *Proceedings of IJCAI*, pages 3311–3317, 2019.
- Karthik Abinav Sankararaman and Aleksandrs Slivkins. Combinatorial semi-bandits with knapsacks. In *Proceedings of AISTATS*, pages 1760–1770, 2018.
- Daniel D Sleator and Robert E Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.
- Aleksandrs Slivkins. Introduction to multi-armed bandits. *Foundations and Trends in Machine Learning*, 12(1-2):1–286, 2019.

## Appendix A. Stochastic Bandits with Knapsacks

In this section we consider the case where the input is sampled i.i.d. from an unknown distribution. We want to still Lagrangify the budget constraint with a parameter  $\lambda$ , but there is a new challenge that we cannot afford losing a multiplicative factor of 2 since we want a  $1 + o(1)$  approximation. The analysis for adversarial arrivals in §A considers two cases, and to balance between them we set  $\lambda = \frac{\text{OPT}_{\text{BWK}}}{2B}$ , but any such analysis can at best only give us  $\text{OPT}_{\text{BWK}}/2$  reward.

For simplicity, let’s first address the above challenge in the  $\ell_\infty$  case. Here the idea, which first appeared in [Badanidiyuru et al. \(2018\)](#), is to introduce a dummy-resource of “time” by adding another dimension. Every action (including the null action) incurs a loss of  $\frac{B}{T}$  in each time step in this new dimension, say dimension 0. Note that since we are working with the  $\ell_\infty$  norm, adding this new dimension does not change  $\text{OPT}_{\text{BWK}}$  since playing  $x^*$  is still feasible. The benefit of introducing this dummy resource is that now our algorithm is guaranteed to exhaust its budget  $B$  by time  $T$ . At the same time, we can assume that  $x^*$  does not exhaust its budget before time  $T$  because

we can scale it down. This implies we are always in Case 1 of the adversarial analysis, and can therefore set  $\lambda \approx \frac{\text{OPT}_{\text{BwK}}}{B}$ .

Although the idea of adding a dummy resource works in the  $\ell_\infty$  case, it's not clear how to implement it for general  $\ell_p$  norms. This is because the dummy resource will contribute to the overall  $\ell_p$  load, unlike in the  $\ell_\infty$  case, which means playing  $x^*$  is no longer feasible and the structure of the problem significantly changes. We still manage to obtain the following sublinear regret.

**Theorem 6** *For stochastic  $\text{BwK}_p$ , there exists an algorithm that given  $\text{OPT}_{\text{BwK}}$  gets a reward*

$$\mathbb{E} \left[ \sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} \right] \geq \text{OPT}_{\text{BwK}} - \text{OPT}_{\text{BwK}} \left( \left( \frac{\|\mathbf{1}_d\|_p}{B} \right)^{1/3} + \left( \frac{p \cdot \|\mathbf{1}_d\|_p}{B} \right)^{1/2} + \frac{\|\mathbf{1}_d\|_p}{B} \text{REGRET} \right),$$

where REGRET is the regret of one-dimensional online learning.

For the important special case of  $\ell_\infty$  norm, we approximate  $\ell_\infty$  by choosing  $p = \frac{\ln d}{\gamma}$ . This gives a bound of  $\mathbb{E} \left[ \sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} \right] \geq \text{OPT}_{\text{BwK}}(1 - 4\gamma)$  whenever  $B \geq \max \left\{ \frac{\ln d \cdot \|\mathbf{1}_d\|_p}{\gamma^3}, \frac{\|\mathbf{1}_d\|_p \cdot \text{REGRET}}{\gamma} \right\}$ .

To prove Theorem 6 (it will follow from Proposition 7), we have to define a different norm.

**A Different Norm.** To overcome the above challenge, we work with a different  $\|\cdot\|_{p,r}$  norm. For any  $(d+1)$ -dimensional vector  $\mathbf{x} = (x_0, \dots, x_d)$ , we define

$$\|\mathbf{x}\|_{p,r} := \left\| \left( x_0, \|(x_1, \dots, x_d)\|_p \right) \right\|_r.$$

The benefit of this norm is that for  $r = \infty$  it allows us to introduce a dummy resource in dimension 0 without affecting the optimum strategy  $x^*$ .

Although powerful, the  $\|\cdot\|_{p,r}$  norm introduces new challenges: We will need to design an approximation  $\Psi_{p,r}(\cdot)$  whose gradient is stable, akin to §2. Fortunately, this is possible by defining

$$\Psi_{p,r}(\mathbf{x}) := \frac{1}{\delta} \left( \left(1 + \delta x_0\right)^r + \left(\|\mathbf{1}_d + \delta \mathbf{y}\|_p\right)^r \right)^{1/r} - \frac{1}{\delta},$$

where  $\delta = \frac{\epsilon}{p+r}$ . In §D we prove the following properties of  $\Psi_{p,r}(\cdot)$ .

**Fact A.1 (Additive Approximation)** *For any integer  $p, r \geq 1$  and load  $\mathbf{\Lambda} \in \mathbb{R}_{\geq 0}^d$ , we have  $\|\mathbf{\Lambda}\|_{p,r} \leq \Psi_{p,r}(\mathbf{\Lambda}) \leq \|\mathbf{\Lambda}\|_{p,r} + \frac{p+r}{\epsilon} \cdot \|\mathbf{1}_d\|_p$ .*

**Fact A.2 (Gradient Stability)** *For any integers  $p, r \geq 1$ , load  $\mathbf{\Lambda} \in \mathbb{R}_{\geq 0}^{d+1}$ , and load increase  $z \in \mathbb{R}_{\geq 0}^{d+1}$  with  $0 \leq z_j \leq 1$  for all  $j$ , coordinate-wise  $\nabla_{p,r} \Psi(\mathbf{\Lambda} + z) \leq (1 + O(\epsilon)) \cdot \nabla_{p,r} \Psi(\mathbf{\Lambda})$ .*

**Fact A.3 (Gradient Norm)** *For any integers  $p, r \geq 1$ , we have  $\|\nabla \Psi_{p,r}\|_{q,s} \leq 1$  where  $q = (1 - 1/p)^{-1}$  and  $1/s = (1 - 1/r)^{-1}$ .*

We now use  $\Psi_{p,r}$  and its properties to define and analyze a surrogate problem.

**Surrogate Online Learning Problem.** We play a surrogate one-dimensional online learning game with  $n$  experts. Besides consuming the  $d$  resources, every action (including null) also consumes  $\frac{B}{T}$  of resource 0. The load  $\mathbf{\Lambda}^{(t)} = \sum_{s=1}^t C^{(s)} \cdot x^{(s)}$  is a  $(d+1)$  dimensional vector where the first row of  $C^{(s)}$  (i.e., dimension 0) contains all  $\frac{B}{T}$ . In time step  $t$  the reward of the  $i$ -th expert is

$$\mathcal{R}_i^{(t)} = \begin{cases} r^{(t)} e_i - \lambda \cdot \langle C^{(t)} e_i, g^{(t)} \rangle & \text{if } \|\mathbf{\Lambda}^{(t-1)}\|_{p,\infty} \leq B \\ 0 & \text{otherwise,} \end{cases}$$

where  $r^{(t)}$  is the reward vector from  $\text{BWK}_p$  problem and  $g^{(t)} := \nabla \Psi_{p,r}(\mathbf{\Lambda}^{(t-1)})$ . We will set

$$\lambda := \frac{\sum_{t=1}^{\tau^*} r^{(t)} \cdot x^*}{B} = \frac{\text{OPT}_{\text{BWK}}}{B} \quad \text{and} \quad \epsilon := \sqrt{\frac{(p+r) \cdot \|\mathbf{1}_d\|_p}{B}}. \quad (8)$$

**Claim A.4** *The reward of each expert per time step is bounded in magnitude by  $(\lambda \cdot \|\mathbf{1}_d\|_p + 1)$ .*

**Proof** Since  $|r^{(t)} e_i| \leq 1$ , we argue  $|\langle C^{(t)} e_i, g^{(t)} \rangle| \leq \|\mathbf{1}\|_p$ . By generalized Cauchy-Schwarz,

$$\langle C^{(t)} e_i, g^{(t)} \rangle \leq \|C^{(t)} e_i\|_{p,r} \cdot \|g^{(t)}\|_{q,s}.$$

This completes the proof because  $\|g^{(t)}\|_{q,s} \leq 1$  by Fact A.3 and  $\|C^{(t)} e_i\|_{p,r} \leq \|\mathbf{1}_d\|_p$ .  $\blacksquare$

Let  $x^*$  be a benchmark solution with the property that  $\|\mathbb{E}[C^{(t)} x^*]\|_{p,r} \leq \frac{B}{T}$ . The following proposition proves Theorem 6 for  $r = \Theta\left(\left(\frac{B}{\|\mathbf{1}_d\|_p}\right)^{1/3}\right)$ .

**Proposition 7** *For stochastic  $\text{BWK}_p$  and any  $r \geq 1$ , the algorithm gets a reward*

$$\mathbb{E} \left[ \sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} \right] \geq \text{OPT}_{\text{BWK}} \left( 1 - \left( \frac{(p+r) \cdot \|\mathbf{1}_d\|_p}{B} \right)^{1/2} - (2^{1/r} - 1) - \frac{\|\mathbf{1}_d\|_p}{B} \cdot \text{REGRET} \right).$$

**Proof** The dummy resource is exhausted after time  $T$ , so the algorithm runs out of some resource before or at time  $T$ . Let  $\tau$  be the random time at which this happens. For analysis, we assume our algorithm stops collecting reward at time  $\tau$ .

We use that  $\text{REGRET}$  bounds  $\sum_{t=1}^T \mathcal{R}^{(t)} \cdot x^{(t)}$  compared to the option of always playing  $x^*$ . By definition  $\mathcal{R}^{(t)} = 0$  for  $t > \tau$ . So, this gives us

$$\sum_{t=1}^{\tau} \mathcal{R}^{(t)} \cdot x^{(t)} \geq \sum_{t=1}^{\tau} \mathcal{R}^{(t)} \cdot x^* - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}. \quad (9)$$

Combining  $\Psi_{p,r}^{(\tau)} \geq \|\mathbf{\Lambda}^{(\tau)}\|_{p,\infty} \geq B$  with Claim 2.4, which holds because of gradient stability (Fact A.2), and using  $(1 + \epsilon)^{-1} \geq (1 - \epsilon)$  gives

$$\sum_{t \leq \tau} \langle C^{(t)} x^{(t)}, g^{(t)} \rangle \geq (1 - \epsilon) \cdot \left( B - \Psi_{p,r}^{(0)} \right).$$

This along with Eq. (9) implies

$$\begin{aligned} \sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} &\geq \lambda \sum_{t=1}^{\tau} \langle C^{(t)} x^{(t)}, g^{(t)} \rangle + \sum_{t=1}^{\tau} \mathcal{R}^{(t)} \cdot x^* - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET} \\ &\geq \lambda (1 - \epsilon) \cdot \left( B - \Psi_{p,r}^{(0)} \right) + \sum_{t=1}^{\tau} \mathcal{R}^{(t)} \cdot x^* - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}. \end{aligned} \quad (10)$$



Now, let us consider  $x^*$ . For  $t \in [T]$ , define a random variable  $Z_t$  with  $Z_t = 1$  if  $\|\mathbf{A}^{(t-1)}\|_{p,\infty} \leq B$  and 0 otherwise. This lets us rewrite  $\mathcal{R}^{(t)} \cdot x^*$  as

$$\mathcal{R}^{(t)} \cdot x^* = Z_t \left( r^{(t)} x^* - \lambda \cdot \langle C^{(t)} x^*, g^{(t)} \rangle \right).$$

Now, we take the expectation over  $r^{(t)}$  and  $C^{(t)}$ , which are by definition independent of  $Z_t$  because  $Z_t$  only depends on steps  $1, \dots, t-1$ . This gives us

$$\mathbb{E}_{r^{(t)}, C^{(t)}}[\mathcal{R}^{(t)} \cdot x^*] = Z_t \left( \mathbb{E}_{r^{(t)}, C^{(t)}}[r^{(t)} x^*] - \lambda \cdot \langle \mathbb{E}_{r^{(t)}, C^{(t)}}[C^{(t)} x^*], g^{(t)} \rangle \right).$$

We have  $\mathbb{E}_{r^{(t)}, C^{(t)}}[r^{(t)} x^*] = \frac{\text{OPT}_{\text{BWK}}}{T}$  and using Hölder's inequality  $\langle \mathbb{E}_{r^{(t)}, C^{(t)}}[C^{(t)} x^*], g^{(t)} \rangle \leq \|\mathbb{E}_{r^{(t)}, C^{(t)}}[C^{(t)} x^*]\|_{p,r} \cdot \|g^{(t)}\|_{q,s} \leq 2^{1/r} \frac{B}{T}$ .

So, in combination after taking the expectation over rounds  $1, \dots, t$ , we have  $\mathbb{E}[\mathcal{R}^{(t)} \cdot x^*] \geq \mathbb{E}[Z_t] \frac{1}{T} \text{OPT}_{\text{BWK}} - \lambda \frac{1}{T} B$  and by linearity of expectation

$$\mathbb{E} \left[ \sum_{t=1}^T \mathcal{R}^{(t)} \cdot x^* \right] \geq \frac{\mathbb{E}[\tau]}{T} \cdot \text{OPT}_{\text{BWK}} - \lambda \cdot 2^{1/r} \cdot \frac{\mathbb{E}[\tau]}{T} B.$$

Combining this with Eq. (10) and by linearity of expectation,

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^{\tau} r^{(t)} \cdot x^{(t)} \right] &\geq \lambda (1 - \epsilon) \cdot \left( B - \Psi_{p,r}^{(0)} \right) + \mathbb{E} \left[ \sum_{t=1}^{\tau} \mathcal{R}^{(t)} \cdot x^* \right] - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET} \\ &\geq \lambda (1 - \epsilon) \cdot \left( B - \Psi_{p,r}^{(0)} \right) + \frac{\mathbb{E}[\tau]}{T} \text{OPT}_{\text{BWK}} - \lambda 2^{1/r} \frac{\mathbb{E}[\tau]}{T} B - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET} \\ &\geq \text{OPT}_{\text{BWK}} \cdot (1 - \epsilon) \cdot \left( 1 - \frac{(p+r) \cdot \|\mathbf{1}_d\|_p}{\epsilon B} \right) - (2^{1/r} - 1) \text{OPT}_{\text{BWK}} - O(\lambda \cdot \|\mathbf{1}_d\|_p) \cdot \text{REGRET}, \end{aligned}$$

where the last inequality uses  $\lambda = \frac{\text{OPT}_{\text{BWK}}}{B}$  and  $\Psi_{p,r}^{(0)} \leq \frac{p+r}{\epsilon} \cdot \|\mathbf{1}_d\|_p$ . Finally, use  $\epsilon$  from Eq. (8). ■

## Appendix B. Lower Bounds for Adversarial Arrivals

In this section we prove our adversarial arrivals lower bounds, both for  $\text{OLVC}_p$  and  $\text{BWK}_p$ . These lower bounds hold even if we assume the algorithm knows the value of  $\text{OPT}_{\text{OLVC}}$  or  $\text{OPT}_{\text{BWK}}$ .

### B.1. Load Balancing

**Proposition 8** *Any online algorithm for adversarial  $\text{OLVC}_p$  is  $\Omega(\min\{p, \log d\})$  competitive.*

**Proof** Our lower bound instance is inspired from the  $\text{OGLB}_p$  lower bound instance of [Awerbuch et al. \(1995\)](#). For simplicity, assume  $d$  is a power of 2. There are  $k = \min\{p, \log_2 d\}$  phases of requests, where the length of each phase is  $L = \frac{T}{k}$ . In the first phase, we think of all the  $d$  dimensions as being *active*. At the end of Phase  $i \in [k]$ , an unbiased random coin  $R_i \in \{0, 1\}$  makes either all the top half (for  $R_i = 0$ ) or the bottom half (for  $R_i = 1$ ) of the currently active dimensions inactive. This means in Phase  $i$  there are exactly  $\frac{d}{2^{i-1}}$  active dimensions. There are  $2^k$  actions where we think of each action  $a$  as a  $k$  bit string. In the  $i$ -th phase, pulling action  $a$  puts a

unit load on all the top or bottom half of active dimensions, depending on whether the  $i$ -th bit of  $a$  is 0 or 1, respectively.

We first observe that for the above instance, the action  $a = (R_1, R_2, \dots, R_p)$  puts exactly  $L$  load in every dimension, which means the  $\ell_p$  norm of its total load is  $L \cdot \|\mathbf{1}\|_p$ . This is true this action puts unit load on exactly those dimensions that will become inactive in the next phase. Thus each arm gets non-zero load in at most one phase.

Next we argue that the  $\ell_p$  norm of the load vector of any online learning algorithm is  $\Omega(L \cdot \|\mathbf{1}\|_p \cdot k)$ . This is because an online algorithm does not know  $R_i$  until the beginning of Phase  $i + 1$ . Thus at best it can evenly distribute between the actions with the  $i$ -th bit being 0/1. Hence its smallest possible  $\ell_p$  norm is given when  $\frac{d}{2}$  dimensions have  $\frac{L}{2}$  load,  $\frac{d}{2^2}$  dimensions have  $\frac{2L}{2}$  load,  $\dots$ ,  $\frac{d}{2^i}$  dimensions have  $\frac{iL}{2}$  load. In particular, each of the  $d/2^k$  dimensions that is active throughout the sequence has a load of  $\frac{Lk}{2}$ . This lower-bounds the  $\ell_p$  norm by  $(\frac{d}{2^k} (\frac{Lk}{2})^p)^{1/p} = \frac{1}{2^{1+k/p}} d^{1/p} Lk = \Theta(L \cdot \|\mathbf{1}\|_p \cdot k)$ .  $\blacksquare$

## B.2. Bandits with Knapsacks

**Proposition 9** *Any online algorithm for adversarial BWK $_p$  is  $\Omega(\min\{p, \log d\})$  competitive.*

**Proof** The idea is to repeat the same instance as the lower bound instance for adversarial OGLB $_p$  with budget  $L \cdot \|\mathbf{1}\|_p$ . Moreover, each non-null action gives a unit-reward. We know from the analysis that there exists an optimal action that satisfies budget for  $k$  phases, resulting in a reward of  $kL$ . Next we argue that every online algorithm playing non-null actions satisfies the budget for at most  $O(1)$  phases.

By repeating the analysis for OLVC $_p$ , we know the best an online algorithm can do is to distribute the load uniformly among the active dimensions. This means if the final reward is  $z$  then all  $d/2^k$  dimensions which are active in all phases must have aggregated a load of  $\frac{z}{2}$  each. This gives a lower-bound for the  $\ell_p$  norm of  $(\frac{d}{2^k} (\frac{z}{2})^p)^{1/p}$ . For feasibility this has to be bounded by budget, i.e.,  $(\frac{d}{2^k} (\frac{z}{2})^p)^{1/p} \leq L \cdot \|\mathbf{1}\|_p = Ld^{1/p}$ , which implies that  $z \leq 2^{1+k/p}L = \Theta(L)$ .  $\blacksquare$

## Appendix C. Removing the Assumption that OPT is Known

For adversarial OLVC $_p$ , it is possible to prove our results from §3 without the assumption that the value of  $\text{OPT}_{\text{OLVC}}$  is known by only losing a small constant factor. For adversarial BWK $_p$ , however, one has to lose an additional  $\Omega(\log T)$  factor when  $\text{OPT}_{\text{BWK}}$  is not known. One cannot avoid this  $\Omega(\log T)$  factor as was shown in [Immorlica et al. \(2019\)](#).

**Load Balancing** For adversarial OLVC $_p$ , we use the standard doubling trick to handle that the algorithm doesn't know the optimal load  $\text{OPT}_{\text{OLVC}}$ . The algorithm operates in phases, where in Phase  $i \geq 1$  it guesses the  $\ell_p$  norm of the optimal solution to be  $2^i$ . Under this assumption, we run the adversarial OLVC $_p$  from §3 where in the beginning of each phase the algorithm imagines that it starts with 0 load in each of the dimensions. If ever during the execution of the algorithm, the load of the algorithm becomes more than  $c \cdot p \cdot 2^i$ , where  $c$  is some sufficiently large constant, the algorithm knows that  $\text{OPT}_{\text{OLVC}}$  must have been higher than  $2^i$  as otherwise it violates the  $O(p)$ -competitive guarantee from §3. Thus the algorithm moves to Phase  $i + 1$ , where its guess for  $\text{OPT}_{\text{OLVC}}$  is  $2^{i+1}$ .

Now we argue that the above algorithm is still  $O(p)$  competitive. This happens because by Minkowski's inequality, we can upper bound the  $\ell_p$  norm of the algorithm's total load as the sum of its  $\ell_p$  load in each phase. Moreover, notice the algorithm never goes past Phase  $i^*$  where  $\text{OPT}_{\text{OLVC}} \in [2^{i^*-1}, 2^{i^*})$ . Hence the total load of the algorithm is at most  $\sum_{i=1}^{i^*} c \cdot p \cdot 2^i = O(p \cdot 2^{i^*}) = O(p \cdot \text{OPT}_{\text{OLVC}})$ .

**Bandits with Knapsacks** For adversarial  $\text{BWK}_p$ , one can get an  $O(\min\{p, \log d\} \log T)$  competitive ratio by first guessing the value of  $\text{OPT}_{\text{BWK}}$  up to a factor of 2. Note that  $\text{OPT}_{\text{BWK}} \in [1, T]$ , so by exponential bucking  $[1, 2), [2, 2^2), \dots, [2^i, 2^{i+1})$ , the algorithm can achieve this with probability at least  $\frac{1}{\log T}$ . Thus the expected reward of the algorithm is  $O(\min\{p, \log d\} \log T)$ -competitive against the fixed optimal action distribution  $x^*$ .

## Appendix D. Missing Proofs

### D.1. Section 2

**Proof** [of Lemma 4] Since  $\Phi$  is twice differentiable with monotone  $\nabla\Phi$ , for any  $\alpha \in \mathbb{R}_{\geq 0}^d$  we have

$$\begin{aligned} \Phi(\mathbf{\Lambda} + z^{(i)}) - \Phi(\mathbf{\Lambda}) &= \int_{s=0}^1 \langle \nabla\Phi(\mathbf{\Lambda} + s \cdot z^{(i)}), z^{(i)} \rangle ds \\ &\leq \int_{s=0}^1 \langle \nabla\Phi(\mathbf{\Lambda} + \alpha + s \cdot z^{(i)}), z^{(i)} \rangle ds = \Phi(\mathbf{\Lambda} + z^{(i)} + \alpha) - \Phi(\mathbf{\Lambda} + \alpha). \end{aligned}$$

Now taking  $\alpha = \sum_{j < i} z^{(j)}$ , we can upper bound

$$\sum_{i=1}^k \left( \Phi(\mathbf{\Lambda} + z^{(i)}) - \Phi(\mathbf{\Lambda}) \right) \leq \sum_{i=1}^k \left( \Phi(\mathbf{\Lambda} + \sum_{j \leq i} z^{(j)}) - \Phi(\mathbf{\Lambda} + \sum_{j < i} z^{(j)}) \right),$$

which equals  $\Phi(\mathbf{\Lambda} + \sum_{i=1}^k z^{(i)}) - \Phi(\mathbf{\Lambda})$  and proves the lemma.  $\blacksquare$

### D.2. Appendix A

**Fact A.2** (*Gradient Stability*) For any integers  $p, r \geq 1$ , load  $\mathbf{\Lambda} \in \mathbb{R}_{\geq 0}^{d+1}$ , and load increase  $z \in \mathbb{R}_{\geq 0}^{d+1}$  with  $0 \leq z_j \leq 1$  for all  $j$ , coordinate-wise  $\nabla_{p,r}\Psi(\mathbf{\Lambda} + z) \leq (1 + O(\epsilon)) \cdot \nabla\Psi_{p,r}(\mathbf{\Lambda})$ .

**Proof** [of Fact A.2] Denote  $\mathbf{y} := (x_1, \dots, x_d)$ . Recall,  $\delta := \frac{\epsilon}{p+r}$  and

$$\Psi_{p,r}(x_0, \mathbf{y}) := \frac{1}{\delta} \left( (1 + \delta x_0)^r + \left( \|\mathbf{1}_d + \delta \mathbf{y}\|_p \right)^r \right)^{1/r} - \frac{1}{\delta}.$$

For ease of notation, denote

$$\Phi_{p,r}(x_0, \mathbf{y}) := \left( 1 + \delta x_0 \right)^r + \left( \|\mathbf{1}_d + \delta \mathbf{y}\|_p \right)^r \quad \text{and} \quad \Phi_p(\mathbf{y}) := \left( \|\mathbf{1}_d + \delta \mathbf{y}\|_p \right)^p.$$

Now we can write the gradient of  $\Psi$  as

$$\frac{\partial \Psi_{p,r}(x_0, \mathbf{y})}{\partial x_0} = \frac{1}{\Phi_{p,r}^{1-1/r}} \left( 1 + \delta x_0 \right)^{r-1} \quad \text{and} \quad \frac{\partial \Psi_{p,r}(x_0, \mathbf{y})}{\partial y_i} = \frac{\Phi_p^{r/p-1}}{\Phi_{p,r}^{1-1/r}} \left( 1 + \delta y_i \right)^{p-1}.$$

We note that if we increase each coordinate of  $(x_0, \mathbf{y})$  by at most 1, then

- the term in the denominators  $\Phi_{p,r}^{1-1/r}$  only increases.
- the terms  $(1 + \delta x_0)^{r-1}$  and  $(1 + \delta y_i)^{p-1}$  increase by a factor at most  $\exp(\epsilon)$  since  $\delta \leq \min\{\frac{1}{p-1}, \frac{1}{r-1}\}$ .
- finally, the term  $\Phi_p^{r/p-1}$  decreases if  $r \leq p$ , and otherwise when  $r > p$  it increases at most by a factor of  $\exp(\delta p \cdot \frac{r}{p}) = 1 + O(\epsilon)$ .

Thus we have that the gradient increases by at most a  $1 + O(\epsilon)$  factor coordinate-wise. ■