



Adversarial Bandits with Knapsacks

NICOLE IMMORLICA, Microsoft Research, USA

KARTHIK SANKARARAMAN, Meta AI, USA

ROBERT SCHAPIRE and ALEKSANDRS SLIVKINS, Microsoft Research, USA

40

We consider *Bandits with Knapsacks* (henceforth, *BwK*), a general model for multi-armed bandits under supply/budget constraints. In particular, a bandit algorithm needs to solve a well-known *knapsack problem*: find an optimal packing of items into a limited-size knapsack. The BwK problem is a common generalization of numerous motivating examples, which range from dynamic pricing to repeated auctions to dynamic ad allocation to network routing and scheduling. While the prior work on BwK focused on the stochastic version, we pioneer the other extreme in which the outcomes can be chosen adversarially. This is a considerably harder problem, compared to both the stochastic version and the “classic” adversarial bandits, in that regret minimization is no longer feasible. Instead, the objective is to minimize the *competitive ratio*: the ratio of the benchmark reward to algorithm’s reward.

We design an algorithm with competitive ratio $O(\log T)$ relative to the best fixed distribution over actions, where T is the time horizon; we also prove a matching lower bound. The key conceptual contribution is a new perspective on the stochastic version of the problem. We suggest a new algorithm for the stochastic version, which builds on the framework of regret minimization in repeated games and admits a substantially simpler analysis compared to prior work. We then analyze this algorithm for the adversarial version, and use it as a subroutine to solve the latter.

Our algorithm is the first “black-box reduction” from bandits to BwK: it takes an arbitrary bandit algorithm and uses it as a subroutine. We use this reduction to derive several extensions.

CCS Concepts: • **Theory of computation** → **Online learning algorithms; Online learning theory; Regret bounds;**

Additional Key Words and Phrases: Multi-armed bandits, adversarial bandits, bandits with knapsacks, regret, competitive ratio, primal-dual algorithms

ACM Reference format:

Nicole Immorlica, Karthik Sankararaman, Robert Schapire, and Aleksandrs Slivkins. 2022. Adversarial Bandits with Knapsacks. *J. ACM* 69, 6, Article 40 (November 2022), 47 pages.

<https://doi.org/10.1145/3557045>

A preliminary extended abstract was published in *FOCS 2019: 60th Annual IEEE Symposium on Foundations of Computer Science*. Some of the results have been strengthened since the conference publications.

Authors’ addresses: N. Immorlica, R. Schapire, and A. Slivkins, Microsoft Research New England, One Memorial Drive, Cambridge, MA, 02142, USA; emails: {nicimm, schapire, slivkins}@microsoft.com; K. Sankararaman, Meta AI, 1 Hacker Way, Menlo Park, CA, 94025; email: karthikabinavs@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0004-5411/2022/11-ART40 \$15.00

<https://doi.org/10.1145/3557045>

1 INTRODUCTION

Multi-armed bandits is a simple abstraction for the tradeoff between *exploration* and *exploitation*, i.e., between making potentially suboptimal decisions for the sake of acquiring new information and using this information for making better decisions. Studied over many decades, multi-armed bandits is a very active research area spanning computer science, operations research, and economics [23, 28, 37, 55, 69, 92].

In this paper, we focus on bandit problems which feature supply or budget constraints, as is the case in many realistic applications. For example, a seller who experiments with prices may have a limited inventory, and a website optimizing ad placement may be constrained by the advertisers' budgets. This general problem is called **Bandits with Knapsacks (BwK)** since, in this model, a bandit algorithm needs effectively to solve a *knapsack problem* (find an optimal packing of items into a limited-size knapsack) or generalization thereof. The BwK model was introduced in [19] as a common generalization of numerous motivating examples, ranging from dynamic pricing to ad allocation to repeated auctions to network routing/scheduling. Various special cases with budget/supply constraints were studied previously, [e.g., 17, 18, 24, 42, 90].

In BwK, the algorithm is endowed with $d \geq 1$ limited resources that are consumed by the algorithm. In each round, the algorithm chooses an action (*arm*) from a fixed set of K actions, and the outcome consists of a reward and consumption of each resource; all are assumed to lie in $[0, 1]$. The algorithm observes *bandit feedback*, i.e., only the outcome of the chosen arm. The algorithm stops at time horizon T , or when the total consumption of some resource exceeds its budget. The goal is to maximize the total reward, denoted REW.

For a concrete example, consider *dynamic pricing*.¹ The algorithm is a seller with limited supply of some product. In each round, a new customer arrives, the algorithm chooses a price, and the customer either buys one item at this price or leaves. A sale at price p implies reward of p and consumption of 1. This example easily extends to $d > 1$ products/resources. Now in each round the algorithm chooses the per-unit price for each resource, and the customer decides how much of each resource to buy at this price.

Prior work on BwK focused on the stochastic version of the problem, called *Stochastic BwK*, where the outcome of each action is drawn from a fixed distribution. This problem has been solved optimally using three different techniques [5, 19], and extended in various directions in subsequent work [5–7, 20].

We go beyond the stochastic version, and instead study the most “pessimistic”, adversarial version where the rewards and resource consumptions can be arbitrary. We call it **adversarial bandits with knapsacks (Adversarial BwK)**, as it extends the classic model of “adversarial bandits” [13]. Bandits aside, this problem subsumes online packing problems [32, 75], where the algorithm observes *full feedback* (the outcomes of all possible actions) in each round, and observes it *before* choosing an action.

Hardness of the problem. Adversarial BwK is a much harder problem compared to Stochastic BwK. The new challenge is that the algorithm needs to decide how much budget to save for the future, without being able to predict it. (It is also the essential challenge in online packing problems, and it drives our lower bounds.) This challenge compounds the ones already present in Stochastic BwK: that exploitation may be severely limited by the resource consumption during exploration, that optimal per-round reward no longer guarantees optimal total reward, and that the best fixed distribution over arms may perform much better than the best fixed arm. Jointly, these challenges amount to the following. An algorithm for Adversarial BwK must compete, during any given time

¹See Section 8 in [19] for a detailed discussion of this and many other examples.

segment $[1, \tau]$, with a distribution over arms that maximizes the total reward on this time segment. However, this distribution may behave very differently, in terms of expected per-round outcomes, compared to the optimal distribution for some other time segment $[1, \tau']$.

In more concrete terms, let OPT_{FD} be the total expected reward of the *best fixed distribution* over arms. In Stochastic BwK (as well as in adversarial bandits) an algorithm can achieve sublinear regret: $\text{OPT}_{\text{FD}} - \mathbb{E}[\text{REW}] = o(T)$.² In contrast, in Adversarial BwK regret minimization is no longer possible, and we therefore are primarily interested in the *competitive ratio* $\text{OPT}_{\text{FD}}/\mathbb{E}[\text{REW}]$.

It is instructive to consider a simple example in which the competitive ratio is at least $\frac{5}{4} - o(1)$ for any algorithm. There are two arms and one resource with budget $\frac{T}{2}$. Arm 1 has zero rewards and zero consumption. Arm 2 has consumption 1 in each round, and offers reward $\frac{1}{2}$ in each round of the first half-time ($\frac{T}{2}$ rounds). In the second half-time, it offers either reward 1 in all rounds, or reward 0 in all rounds. Thus, there are two problem instances that coincide for the first half-time and differ in the second half-time. The algorithm needs to choose how much budget to invest in the first half-time, without knowing what comes in the second. Any choice leads to competitive ratio at least $\frac{5}{4}$ on one of the problem instances.

Extending this idea, we prove an even stronger lower bound on the competitive ratio:

$$\text{OPT}_{\text{FD}}/\mathbb{E}[\text{REW}] \geq \Omega(\log T). \quad (1.1)$$

Like the simple example above, the lower-bounding construction involves only two arms and only one resource, and forces the algorithm to make a huge commitment without knowing the future.

Algorithmic contributions. Our main result is an algorithm which nearly matches (1.1), achieving

$$\mathbb{E}[\text{REW}] \geq \frac{1}{O(\log T)} (\text{OPT}_{\text{FD}} - \text{reg}), \quad (1.2)$$

where reg is a low-order regret term.

We put forward a new algorithm for BwK, called LagrangeBwK, that unifies the stochastic and adversarial versions. It has a natural game-theoretic interpretation for Stochastic BwK, and admits a simpler analysis compared to the prior work. For Adversarial BwK, we use LagrangeBwK as a subroutine, though with a different parameter and a different analysis, to derive two algorithms: a simple one that achieves (1.2), and a more involved one that achieves the same competitive ratio with high probability. Absent resource consumption, we recover the optimal $\tilde{O}(\sqrt{KT})$ regret for adversarial bandits.

LagrangeBwK is based on a new perspective on Stochastic BwK. We reframe a standard linear relaxation for Stochastic BwK in a way that gives rise to a repeated zero-sum game, where the two players choose among arms and resources, respectively, and the payoffs are given by the Lagrange function of the linear relaxation. Our algorithm consists of two online learning algorithms playing this repeated game. We analyze LagrangeBwK for Stochastic BwK, building on the tools from regret minimization in stochastic games, and achieve a near-optimal regret bound when the optimal value and the budgets are $\Omega(T)$.³

We obtain several extensions, where we derive improved performance guarantees for some scenarios. These extensions showcase the *modularity* of LagrangeBwK, in the sense that the two players can be implemented as arbitrary algorithms for adversarial online learning that admit a given regret bound. Each extension follows from the main results, with a different choice of the players' algorithms.

²More specifically, one can achieve regret $\tilde{O}(\sqrt{KT})$ for adversarial bandits [13], as well as for Stochastic BwK if all budgets are $\Omega(T)$ [19]. One can achieve sublinear regret for Stochastic BwK if all budgets are $\Omega(T^\alpha)$, $\alpha \in (0, 1)$ [19].

³This regime is of primary importance in prior work [e.g., 24, 99].

Discussion. LagrangeBwK has numerous favorable properties. As just discussed, it is simple, unifying, modular, and yields strong performance guarantees in multiple settings. It is the first “black-box reduction” from bandits to BwK: we take a bandit algorithm and use it as a subroutine for BwK. This is a very natural algorithm for the stochastic version once the single-shot game is set up; indeed, it is immediate from prior work that the repeated game converges to the optimal distribution over arms. Its regret analysis for Stochastic BwK is extremely clean. Compared to prior work [5, 19], we side-step the intricate analysis of sensitivity of the linear program to non-uniform stochastic deviations that arise from adaptive exploration.

LagrangeBwK has a primal-dual interpretation, as arms and resources correspond respectively to primal and dual variables in the linear relaxation. Two players in the repeated game can be seen as the respective *primal algorithm* and *dual algorithm*. Compared to the rich literature on *primal-dual algorithms* [32, 75, 101] [including the more recent literature on stochastic online packing problems 8, 44, 45, 48, 77] LagrangeBwK has a specific and modular structure dictated by the repeated game.

Logarithmic competitive ratios are fairly common and well-accepted in the area of approximation algorithms, and particularly in online algorithms (see Related Work for citations).

Benchmarks. We argue that the best fixed distribution over arms is an appropriate benchmark for Adversarial BwK. First, consider the total expected reward of the *best dynamic policy*, denote it OPT_{DP} . (The best dynamic policy is the best algorithm, in hindsight, that is allowed to switch arms arbitrarily across time-steps.) This is the strongest possible benchmark, but it is *too* strong for Adversarial BwK. Indeed, we show a simple example with just one resource (with budget B), where competitive ratio against this benchmark is at least $\frac{T}{B^2}$ for any algorithm. Second, consider the total expected reward of the *best fixed arm*, denote it OPT_{FA} . It is a traditional benchmark in multi-armed bandits, but is uninteresting for Adversarial BwK. We show that the competitive ratio is at least $\Omega(K)$ in the worst case, and this is matched, in expectation and up to a constant factor, by a trivial algorithm that samples one arm at random and sticks with it forever.

For Stochastic BwK, these three benchmarks are related as follows. The best fixed distribution is still the main object of interest in the analysis. However, all results – both ours and prior work – are almost automatically extended to compete against the best dynamic policy. The best fixed arm is a much weaker benchmark than the best fixed distribution: there are simple examples when their expected reward differs by a factor of two, in multiple special cases of interest [19].

2 RELATED WORK

The literature on regret-minimizing online learning is vast; see [28, 37, 60] for background. Most relevant are two algorithms for adversarial rewards/costs: Hedge for full feedback [53], and EXP3 for bandit feedback [13]; both are based on the weighted majority algorithm from [70].

Stochastic BwK was introduced and optimally solved in [19]. Subsequent work extended these results to soft supply/budget constraints [5], a more general notion of rewards⁴ [5], combinatorial semi-bandits [86], and contextual bandits [6, 7, 20]. Several special cases with budget/supply constraints were studied previously: dynamic pricing [17, 24, 25, 99], dynamic procurement [18, 90] (a version of dynamic pricing where the algorithm is a buyer rather than a seller), dynamic ad allocation [42, 91], and a version with a single resource and unlimited time [46, 58, 95, 96]. While all this work is on regret minimization, Guha and Munagala [56], Gupta et al. [57] studied closely related Bayesian formulations.

Stochastic BwK was optimally solved using three different algorithms [5, 19], with extremely technical and delicate analyses. All three algorithms involve inherently ‘stochastic’ techniques

⁴The total reward is determined by the time-averaged outcome vector, but can be an arbitrary Lipschitz-concave function thereof.

such as “successive elimination” and “optimism under uncertainty”, and do not appear to extend to the adversarial version. One of them, PrimalDualBwK from [19], is a primal-dual algorithm superficially similar to ours. Indeed, it decouples into two online learning algorithms: a “primal” algorithm which chooses among arms, and a “dual” algorithm similar to ours, which chooses among resources. However, the two algorithms are not playing a repeated game in any meaningful sense, let alone a zero-sum game. The primal algorithm operates under a much richer input: it takes the entire outcome vector for the chosen arm, as well as the “dual distribution” – the distribution over resources chosen by the dual algorithm. Further, the primal algorithm is very problem-specific: it interprets the dual distribution as a vector of costs over resources, and chooses arms with the largest reward-to-cost ratios, estimated using “optimism under uncertainty”.

Our approach to using regret minimization in games can be traced to [51, 53] (see Ch. 6 in [87]), who showed how a repeated zero-sum game played by two agents yields an approximate Nash equilibrium. This approach has been used as a unifying algorithmic framework for several problems: boosting [51], linear programs [10], maximum flow [41], and convex optimization [1, 98]. While we use a result with the $1/\sqrt{t}$ convergence rate for the equilibrium property, recent literature obtains faster convergence for cumulative payoffs (but not for the equilibrium property) under various assumptions [43, 80, 100].

Repeated Lagrangian games, in conjunction with regret minimization in games, have been used in a series of recent papers [2, 62, 65, 83–85], as an algorithmic tool to solve convex optimization problems; application domains range from differential privacy to algorithmic fairness to learning from revealed preferences. All these papers deal with deterministic games (i.e., same game matrix in all rounds). Reframing the problem in terms of repeated Lagrangian games is a key technical insight in this work. Most related to our paper are Roth et al. [84, 85], where a repeated Lagrangian game is used as a subroutine (the “inner loop”) in an online algorithm; the other papers solve an offline problem. We depart from this prior work in several respects. Our main results are for the adversarial version, where the standard machinery does not apply and we provide a very different analysis. For the stochastic version, we use a stochastic game and we deal with some subtleties specific to BwK.

Online packing problems [e.g., 32, 33, 45] can be seen as a special case of Adversarial BwK with a much more permissive feedback model: the algorithm observes full feedback (the outcomes for all actions) before choosing an action. Online packing subsumes various *online matching* problems, including the *AdWords problem* [76] motivated by ad allocation [see 75, for a survey]. While we derive $O(\log T)$ competitive ratio against OPT_{FD} , online packing admits a similar result against OPT_{DP} .

Another related line of work concerns online convex optimization with constraints [39, 40, 73, 74, 78]. Their setting differs from ours in several important respects. First, the action set is a convex subset of \mathbb{R}^K (and the algorithms rely on the power to choose arbitrary actions in this set). In particular, there is no immediate way to handle discrete action sets.⁵ Second, convexity/concavity is assumed on the rewards and resource consumption. Third, full feedback is observed for the resource consumption. Moreover, in all papers except Chen and Giannakis [39], one also observes either full feedback on rewards or the rewards gradient around the chosen action. Fourth, their algorithm only needs to satisfy the budget constraints at the time horizon (whereas in BwK the budget constraints hold for all rounds). Fifth, their fixed-distribution benchmark is weaker than ours: essentially, its time-averaged consumption must be small enough at each round t . Due to these differences, this setting admits sublinear regret for adversarial outcomes [78]. The other papers in this line of work focus on stochastic outcomes.

⁵Unless there is full feedback, in which case one can use a standard reduction whereby actions in online convex optimization correspond to distributions over actions in a K -armed bandit problem.

Logarithmic competitive ratios are quite common in prior work on approximation algorithms and online algorithms. Examples include: set cover [63, 71], buy-at-bulk network design [15], sparsest cut [11], the dial-a-ride problem [38], online k-server [21], online packing/covering [16], online set cover [9], online network design [97], and online paging [49].

2.1 Simultaneous and Independent Work

Three related papers have come to our attention after the initial version of our paper appeared on arxiv.org in Nov'18. At the time, Cardoso et al. [35], Rangi et al. [82] were available as yet unpublished technical reports, and Cardoso et al. [34] had not yet appeared.

Cardoso et al. [35] consider online convex optimization with knapsacks: essentially, the problem defined in Section 7.4, but with full feedback. Focusing on the stochastic version, they design an algorithm similar to LagrangeBwK, and derive a regret bound similar to ours, using a similar analysis. They also claim an extension to bandit feedback, without providing any details (such as the precise statement of Lemma 3.1 in terms of the regret property (3.2)).

Rangi et al. [82] consider Adversarial BwK in the special case when there is only one constrained resource, including time. They attain sublinear regret, i.e., a regret bound that is sublinear in T . They also assume a known lower bound $c_{\min} > 0$ on realized per-round consumption of each resource, and their regret bound scales as $1/c_{\min}$. They also achieve $\text{polylog}(T)$ instance-dependent regret for the stochastic version using the same algorithm (matching results from prior work on the stochastic version). BwK with only one constrained resource (including time) is a much easier problem, compared to the general case with multiple resources studied in this paper, in the following sense. First, the single-resource version admits much stronger performance guarantees ($\text{polylog}(T)$ vs. \sqrt{T} regret bounds for Stochastic BwK, and sublinear regret vs. competitive ratio for Adversarial BwK). Second, the optimal all-knowing time-invariant policy is the best fixed arm, rather than the best fixed distribution over arms.

Cardoso et al. [34] study online learning in repeated adversarial zero-sum games (which is our main technical tool). They obtain a powerful result for arbitrary games: an online learning algorithm which controls both players and guarantees convergence to the Nash equilibrium. They apply their framework to train **Generative Adversarial Networks (GANs)**. Interestingly, they achieve the competitive ratio of 1, despite the adversarial setting. Their algorithm can continue up to round T , with no stopping rule like in BwK; for this reason, their results do not have an immediate bearing on our problem.

3 PRELIMINARIES

We use bold fonts to represent vectors and matrices. We use standard notation whereby, for a positive integer K , $[K]$ stands for $\{1, 2, \dots, K\}$, and Δ_K denotes the set of all probability distributions on $[K]$. Some of the notation introduced further is summarized in Appendix B.

Bandits with Knapsacks (BwK). There are T rounds, K possible actions and d resources, indexed as $[T], [K], [d]$, respectively. In each round $t \in [T]$, the algorithm chooses an action $a_t \in [K]$ and receives an *outcome vector* $\mathbf{o}_t = (r_t; c_{t,1}, \dots, c_{t,d}) \in [0, 1]^{d+1}$, where r_t is a reward and $c_{t,i}$ is consumption of each resource $i \in [d]$. Each resource i is endowed with budget $B_i \leq T$. The game stops early, at some round $\tau_{\text{alg}} < T$, when/if the total consumption of any resource exceeds its budget. The algorithm's objective is to maximize its total reward. Without loss of generality all budgets are the same: $B_1 = B_2 = \dots = B_d = B$.⁶

⁶To see that this is indeed w.l.o.g., for each resource i , divide all per-round consumptions $c_{t,i}$ by B_i/B , where $B := \min_{i \in [d]} B_i$ is the smallest budget. In the modified problem instance, all consumptions still lie in $[0, 1]$, and all the budgets are equal to B .

Given: action set A , payoff range $[b_{\min}, b_{\max}]$.

In each round $t \in [T]$,

1. the adversary chooses a payoff vector $\mathbf{f}_t \in [b_{\min}, b_{\max}]^K$;
2. the algorithm chooses a distribution \mathbf{p}_t over A , without observing \mathbf{f}_t ;
3. algorithm's chosen action $a_t \in A$ is drawn independently from \mathbf{p}_t ;
4. payoff $f_t(a_t)$ is received by the algorithm.

Fig. 1. Adversarial online learning.

The outcome vectors are chosen as follows. In each round t , the adversary chooses the *outcome matrix* $\mathbf{M}_t \in [0, 1]^{K \times (d+1)}$, where rows correspond to actions. The outcome vector \mathbf{o}_t is defined as the a_t -th row of this matrix, denoted $\mathbf{M}_t(a_t)$. Only this row is revealed to the algorithm. The adversary is deterministic and *oblivious*, meaning that the entire sequence $\mathbf{M}_1, \dots, \mathbf{M}_T$ is chosen before round 1. A problem instance of BwK consists of (known) parameters (d, K, T, B) , and the (unknown) sequence $\mathbf{M}_1, \dots, \mathbf{M}_T$.

In the **stochastic version of BwK**, henceforth termed **Stochastic BwK**, each outcome matrix \mathbf{M}_t is chosen from some fixed but unknown distribution \mathcal{D}_{BwK} over the outcome matrices. An instance of this problem consists of (known) parameters (d, K, T, B) , and the (unknown) distribution \mathcal{D}_{BwK} .

Following prior work [5, 19], we assume, w.l.o.g., that one of the resources is a *dummy resource* similar to time; formally, each action consumes B/T units of this resource per round (we only need this for Stochastic BwK). Further, we posit that one of the actions is a *null action*, which lets the algorithm skip a round: it has 0 reward and consumes 0 amount of each resource other than the dummy resource.

Benchmarks. Let $\text{REW}(\text{ALG}) = \sum_{t \in [\tau_{\text{alg}}]} r_t$ be the total reward of algorithm ALG in the BwK problem. Our benchmark is the *best fixed distribution*, a distribution over actions which maximizes $\mathbb{E}[\text{REW}(\cdot)]$ for a particular problem instance. The expected total reward of this distribution is denoted OPT_{FD} .

For Stochastic BwK, one can compete with the *best dynamic policy*: an algorithm that maximizes $\mathbb{E}[\text{REW}(\cdot)]$ for a particular problem instance. Essentially, this algorithm knows the latent distribution \mathcal{D}_{BwK} over outcome matrices. Its expected total reward is denoted OPT_{DP} .

Adversarial online learning. To state the framework of “regret minimization in games” below, we need to introduce the protocol of *adversarial online learning*, see Figure 1.

In this protocol, the adversary can use previously chosen arms to choose the payoff vector \mathbf{f}_t , but not the algorithm's random seed. The distribution \mathbf{f}_t is chosen as a deterministic function of history. (The history at round t consists, for each round $s < t$, of the chosen action a_s and the observed feedback in this round.) We focus on two feedback models: *bandit feedback* (no auxiliary feedback) and *full feedback* (the entire payoff vector \mathbf{f}_t). The version for costs can be defined similarly, by setting the payoffs to be the negative of costs.

We are interested in adversarial online learning algorithms with known upper bounds on *regret*,

$$R_{\text{AOL}}(T) := \left[\max_{a \in A} \sum_{t \in [T]} f_t(a) \right] - \left[\sum_{t \in [T]} f_t(a_t) \right]. \quad (3.1)$$

The benchmark here is the total payoff of the best arm, according to the payoff vectors actually chosen by the adversary. More precisely, we assume high-probability regret bounds of the following

form:

$$\forall \delta > 0 \quad \Pr [R_{\text{AOL}}(T) \leq (b_{\max} - b_{\min}) R_{\delta}(T)] \geq 1 - \delta, \quad (3.2)$$

for some function $R_{\delta}(\cdot)$. We will actually use a stronger version implied by (3.2),⁷

$$\forall \delta > 0 \quad \Pr [\forall \tau \in [T] \quad R_{\text{AOL}}(\tau) \leq (b_{\max} - b_{\min}) R_{\delta/T}(T)] \geq 1 - \delta. \quad (3.3)$$

Algorithms EXP3.P [13] for bandit feedback, and Hedge [52] for full feedback, satisfy (3.2) with, respectively,

$$R_{\delta}(T) = O\left(\sqrt{|A| T \log(T/\delta)}\right) \text{ and } R_{\delta}(T) = O\left(\sqrt{T \log(|A|/\delta)}\right). \quad (3.4)$$

Regret minimization in games. We build on the framework of *regret minimization in games*. A *zero-sum game* (A_1, A_2, G) is a game between two players $i \in \{1, 2\}$ with action sets A_1 and A_2 and payoff matrix $G \in \mathbb{R}^{A_1 \times A_2}$. If each player i chooses an action $a_i \in A_i$, the outcome is a number $G(a_1, a_2)$. Player 1 receives this number as *reward*, and player 2 receives it as *cost*. A *repeated zero-sum game* \mathcal{G} with action sets A_1 and A_2 , time horizon T and game matrices $G_1, \dots, G_T \in \mathbb{R}^{A_1 \times A_2}$ is a game between two algorithms, ALG_1 and ALG_2 , which proceeds over T rounds such that each round t is a zero-sum game (A_1, A_2, G_t) . The goal of ALG_1 is to maximize the total reward, and the goal of ALG_2 is to minimize the total cost.

The game \mathcal{G} is called *stochastic* if the game matrix G_t in each round t is drawn independently from some fixed distribution. For such games, we are interested in the *expected game*, defined by the expected game matrix $G = \mathbb{E}[G_t]$. We can relate the algorithms' performance to the minimax value of G .

LEMMA 3.1. *Consider a stochastic repeated zero-sum game between algorithms ALG_1 and ALG_2 , with payoff range $[b_{\min}, b_{\max}]$. Assume that each ALG_j , $j \in \{1, 2\}$ is an algorithm for adversarial online learning, as per Figure 1, which satisfies regret bound (3.2) with $R_{\delta}(T) = R_{j,\delta}(T)$.*

Let τ be some fixed round in the game. For each algorithm ALG_j , $j \in \{1, 2\}$, let A_j be its action set, let $p_{t,j} \in \Delta_{A_j}$ be the distribution chosen in each round t , and let $\bar{p}_j = \frac{1}{\tau} \sum_{t \in [\tau]} p_{t,j}$ be the average play distribution at round τ . Let v^ be the minimax value for the expected game $G = \mathbb{E}[G_t]$.*

Then for each $\delta > 0$, with probability at least $1 - 2\delta$ it holds that

$$\forall p_2 \in \Delta_{A_2} \quad \bar{p}_1^T G p_2 \geq v^* - \frac{1}{\tau} (b_{\max} - b_{\min}) \left(R_{1,\delta/T}(T) + R_{2,\delta/T}(T) + 4\sqrt{2T \log(T/\delta)} \right). \quad (3.5)$$

Equation (3.5) states that the average play of player 1 is approximately optimal against any distribution chosen by player 2.⁸ This lemma is well-known for the deterministic case (i.e., when $G_t = G$ for each round t), and folklore for the stochastic case. We provide a proof in Appendix A.4 for the sake of completeness.

4 A NEW ALGORITHM FOR STOCHASTIC BWK

We present a new algorithm for Stochastic BwK, based on the framework of regret minimization in games. This is a very natural algorithm once the single-shot game is set up, and it allows for a very clean regret analysis. We will also use this algorithm as a subroutine for the adversarial version.

On a high level, we define a stochastic zero-sum game for which a mixed Nash equilibrium corresponds to an optimal solution for a linear relaxation of the original problem. Our algorithm

⁷Regret bound (3.3) follows from (3.2) using a simple “zeroing-out” trick: for a given round $\tau \in [T]$, the adversary can set all future payoffs to some fixed value $x \in [b_{\min}, b_{\max}]$, in which case $R_{\text{AOL}}(\tau) = R_{\text{AOL}}(T)$.

⁸If each player j chooses distribution $p_j \in \Delta_{A_j}$, and the game matrix is G , then expected reward/cost is $\bar{p}_1^T G p_2$.

consists of two regret-minimizing algorithms playing this game. The framework of regret minimization in games guarantees that the average primal and dual play distributions (\bar{p}_1 and \bar{p}_2 in Lemma 3.1) approximate the mixed Nash equilibrium in the expected game, which correspondingly approximates the optimal solution.

4.1 Linear Relaxation and Lagrange Functions

We start with a linear relaxation of the problem that all prior work relies on. This relaxation is stated in terms of expected rewards/consumptions, i.e., implicitly, in terms of the expected outcome matrix $\mathbf{M} = \mathbb{E}[\mathbf{M}_t]$. We explicitly formulate the relaxation in terms of \mathbf{M} , and this is essential for the subsequent developments. For ease of notation, we write the a -th row of \mathbf{M} , for each action $a \in [K]$, as

$$\mathbf{M}(a) = (r^{\mathbf{M}}(a); c_1^{\mathbf{M}}(a), \dots, c_d^{\mathbf{M}}(a)),$$

so that $r^{\mathbf{M}}(a)$ is the expected reward and $c_i^{\mathbf{M}}(a)$ is the expected consumption of each resource i .

Essentially, the relaxation assumes that each instantaneous outcome matrix \mathbf{M}_t is equal to the expected outcome matrix $\mathbf{M} = \mathbb{E}[\mathbf{M}_t]$. The relaxation seeks the best distribution over actions, focusing on a single round with budgets rescaled as B/T . Thus, we have the following **linear program (LP)**:

$$\begin{aligned} & \text{maximize} && \sum_{a \in [K]} X(a) r^{\mathbf{M}}(a) && \text{such that} \\ & && \sum_{a \in [K]} X(a) = 1 \\ & \forall i \in [d] && \sum_{a \in [K]} X(a) c_i^{\mathbf{M}}(a) \leq B/T \\ & \forall a \in [K] && 0 \leq X(a) \leq 1. \end{aligned} \quad (4.1)$$

We denote this LP by $\text{LP}_{\mathbf{M}, B, T}$. The solution \mathbf{X} is the best fixed distribution over actions, according to the relaxation. The value of this LP, denoted $\text{OPT}_{\text{LP}}(\mathbf{M}, B, T)$, is the expected per-round reward of this distribution. It is also the total reward of \mathbf{X} in the relaxation, divided by T . We know from [19] that

$$T \cdot \text{OPT}_{\text{LP}}(\mathbf{M}, B, T) \geq \text{OPT}_{\text{DP}} \geq \text{OPT}_{\text{FD}}, \quad (4.2)$$

where OPT_{DP} and OPT_{FD} are the total expected rewards of, respectively, the best dynamic policy and the best fixed distribution. In words, OPT_{DP} is sandwiched between the total expected reward of the best fixed distribution and that of its linear relaxation.

Associated with the linear program $\text{LP}_{\mathbf{M}, B, T}$ is the *Lagrange function* $\mathcal{L} = \mathcal{L}_{\mathbf{M}, B, T}$. It is a function $\mathcal{L} : \Delta_K \times \mathbb{R}_{\geq 0}^d \rightarrow \mathbb{R}$ defined as

$$\mathcal{L}(\mathbf{X}, \boldsymbol{\lambda}) := \sum_{a \in [K]} X(a) r^{\mathbf{M}}(a) + \sum_{i \in [d]} \lambda_i \left[1 - \frac{T}{B} \sum_{a \in [K]} X(a) c_i^{\mathbf{M}}(a) \right]. \quad (4.3)$$

The values $\lambda_1, \dots, \lambda_d$ in Equation (4.3) are called the *dual variables*, as they correspond to the variables in the dual LP. Lagrange functions are meaningful due to their max-min property (e.g., Theorem D.2.2 in [22]):

$$\min_{\lambda \geq 0} \max_{\mathbf{X} \in \Delta_K} \mathcal{L}(\mathbf{X}, \boldsymbol{\lambda}) = \max_{\mathbf{X} \in \Delta_K} \min_{\lambda \geq 0} \mathcal{L}(\mathbf{X}, \boldsymbol{\lambda}) = \text{OPT}_{\text{LP}}(\mathbf{M}, B, T). \quad (4.4)$$

This property holds for our setting because $\text{LP}_{\mathbf{M}, B, T}$ has at least one feasible solution (namely, one that puts probability one on the null action), and the optimal value of the LP is bounded.

Remark 4.1. We use the linear program $\text{LP}_{\mathbf{M}, B, T}$ and the associated Lagrange function $\mathcal{L}_{\mathbf{M}, B, T}$ throughout the paper. Both are parameterized by an outcome matrix \mathbf{M} , budget B and time horizon T . In particular, we can plug in an arbitrary \mathbf{M} , and we heavily use this ability throughout. For the adversarial version, it is essential to plug in parameter $T_0 \leq T$ instead of the time horizon T . For

the analysis of the high-probability result in Adversarial BwK, we use a rescaled budget $B_0 \leq B$ instead of budget B .

4.2 Our Algorithm: Repeated Lagrangian Game

The Lagrange function $\mathcal{L} = \mathcal{L}_{M,B,T}$ from (4.3) defines the following zero-sum game: the *primal player* chooses an arm a , the *dual player* chooses a resource i , and the payoff is a number

$$\mathcal{L}(a, i) = r^M(a) + 1 - \frac{T}{B} c_i^M(a). \quad (4.5)$$

The primal player receives this number as a reward, and the dual player receives it as cost. This game is termed the *Lagrangian game* induced by $\mathcal{L}_{M,B,T}$. This game will be crucial throughout the paper.

The Lagrangian game is related to the original linear program as follows:

LEMMA 4.2. *Assume one of the resources is the dummy resource. Consider the linear program $\text{LP}_{M,B,T}$, for some outcome matrix M . Then the value of this LP equals the minimax value v^* of the Lagrangian game induced by $\mathcal{L}_{M,B,T}$. Further, if (X, λ) is a mixed Nash equilibrium in the Lagrangian game, then X is an optimal solution to the LP.*

The proof can be found in Appendix A.2. The idea is that because of the special structure of the LP, the second equality in (4.4) also holds when the dual vector λ is restricted to distributions.

Consider a repeated version of the Lagrangian game. Formally, the *repeated Lagrangian game* with parameters $B_0 \leq B$ and $T_0 \leq T$ is a repeated zero-sum game between the *primal algorithm* that chooses among arms and the *dual algorithm* that chooses among resources. Each round t of this game is the Lagrangian game induced by the Lagrange function $\mathcal{L}_t := \mathcal{L}_{M_t, B_0, T_0}$, where M_t is the round- t outcome matrix. Note that we use parameters B_0, T_0 instead of budget B and time horizon T .⁹

Remark 4.3. Consider repeated Lagrangian game for Stochastic BwK (with $B_0 = B$ and $T_0 = T$). The payoffs in the expected game are defined by the expected Lagrange function $\mathcal{L} := \mathbb{E}[\mathcal{L}_t]$. By linearity, \mathcal{L} is the Lagrange function for the expected outcome matrix $M = \mathbb{E}[M_t]$:

$$\mathcal{L} := \mathbb{E}[\mathcal{L}_t] = \mathcal{L}_{M, B, T}. \quad (4.6)$$

Our algorithm, called LagrangeBwK, is very simple: it is a repeated Lagrangian game in which the primal algorithm receives bandit feedback, and the dual algorithm receives full feedback.

Let a_t and i_t be, respectively, the chosen arm and resource in round t . The payoff is therefore $\mathcal{L}_t(a_t, i_t)$. It can be rewritten in terms of the observed outcome vector $\mathbf{o}_t = (r_t; c_{t,1}, \dots, c_{t,d})$ (which corresponds to the a_t -th row of the instantaneous outcome matrix M_t):

$$\mathcal{L}_t(a_t, i_t) = r_t + 1 - \frac{T_0}{B_0} c_{t,i_t} \in \left[-\frac{T_0}{B_0} + 1, 2 \right]. \quad (4.7)$$

Note that the payoff range is $[b_{\min}, b_{\max}] = \left[-\frac{T_0}{B_0} + 1, 2 \right]$.

With this notation, the pseudocode for LagrangeBwK is summarized in Algorithm 1. The pseudocode is simple and self-contained, without referring to the formalism of repeated games and Lagrangian functions. Note that the algorithm is implementable, in the sense that the outcome vector \mathbf{o}_t revealed in each round t of the BwK problem suffices to generate full feedback for the dual algorithm.

⁹These parameters are needed only for the adversarial version. For Stochastic BwK we use $B_0 = B$ and $T_0 = T$.

ALGORITHM 1: Algorithm LagrangeBwK for Stochastic BwK.

input: parameters B_0, T_0 , primal algorithm ALG_1 , dual algorithm ALG_2 .
 // $\text{ALG}_1, \text{ALG}_2$ are adversarial online learning algorithms
 // with bandit feedback and full feedback, respectively
for round $t = 1, 2, 3, \dots$ **do**

(1) ALG_1 returns arm $a_t \in [K]$, algorithm ALG_2 returns resource $i_t \in [d]$.
 (2) arm a_t is chosen,
 outcome vector $\mathbf{o}_t = (r_t(a_t); c_{t,1}(a_t), \dots, c_{t,d}(a_t)) \in [0, 1]^{d+1}$ is observed.
 (3) The payoff $\mathcal{L}_t(a_t, i_t)$ from (4.7) is reported to ALG_1 as reward, and to ALG_2 as cost.
 (4) The payoff $\mathcal{L}_t(a_t, i)$ is reported to ALG_2 for each resource $i \in [d]$.

4.3 Performance Guarantees

We consider algorithm LagrangeBwK with parameter $T_0 = T$. We assume the existence of the dummy resource; this is to ensure that the crucial step, Equation (4.13), works out even if the algorithm stops at time T , without exhausting any actual resources. We obtain a regret bound that is non-trivial whenever $B > \Omega(\sqrt{T})$, and is optimal, up to log factors, in the regime when $\min(\text{OPT}_{\text{DP}}, B) > \Omega(T)$.

THEOREM 4.4. *Consider Stochastic BwK with K arms, d resources, time horizon T , and budget B . Assume that one resource is the dummy resource (with consumption $\frac{B}{T}$ for each arm). Fix the failure probability parameter $\delta \in (0, 1)$. Consider algorithm LagrangeBwK with parameters $B_0 = B, T_0 = T$.*

If EXP3.P and Hedge are used as the primal and the dual algorithms, respectively, then the algorithm achieves the following regret bound, with probability at least $1 - \delta$:

$$\text{OPT}_{\text{DP}} - \text{REW}(\text{LagrangeBwK}) \leq O\left(\frac{T}{B} \sqrt{TK \log(dT/\delta)}\right). \quad (4.8)$$

In general, suppose each algorithm ALG_j satisfies a regret bound (3.2) with $R_\delta(T) = R_{j,\delta}(T)$ and payoff range $[b_{\min}, b_{\max}] = [-\frac{T}{B} + 1, 2]$. Then with probability at least $1 - O(\delta T)$ it holds that

$$\text{OPT}_{\text{DP}} - \text{REW}(\text{LagrangeBwK}) \leq O\left(\frac{T}{B}\right) \left(R_{1,\delta/T}(T) + R_{2,\delta/T}(T) + \sqrt{T \log(dT/\delta)}\right). \quad (4.9)$$

Remark 4.5. To obtain (4.8) from the “black-box” result (4.9), we use regret bounds in Equation (3.4).

Remark 4.6. From [19], the optimal regret bound for Stochastic BwK is

$$\text{OPT}_{\text{DP}} - \mathbb{E}[\text{REW}] \leq \tilde{O}\left(\sqrt{K \text{OPT}_{\text{DP}}}\left(1 + \sqrt{\text{OPT}_{\text{DP}}/B}\right)\right).$$

Thus, the regret bound (4.8) is near-optimal if $\min(\text{OPT}_{\text{DP}}, B) > \Omega(T)$, and non-trivial if $B > \Omega(\sqrt{T})$.

We next prove the “black-box” regret bound (4.9). For the sake of analysis, consider a version of the repeated Lagrangian game that continues up to the time horizon T . In what follows, we separate the “easy steps” from what we believe is the crux of the proof.

Notation. Let X_t be the distribution chosen in round t by the primal algorithm ALG_1 . Let $\bar{X}_\tau := \frac{1}{\tau} \sum_{t \in [\tau]} X_t$ be the distribution of average play up to round τ . Let $\mathbf{M} = \mathbb{E}[\mathbf{M}_t]$ be the expected outcome matrix. Let $\mathbf{r} = (r^{\mathbf{M}}(a) : a \in [K])$ be the vector of expected rewards over the actions. Likewise, $\mathbf{c}_i = (c_i^{\mathbf{M}}(a) : a \in [K])$ be the vector of expected consumption of each resource $i \in [d]$.

Using Azuma-Hoeffding inequality. Consider the first τ rounds, for some $\tau \in [T]$. The average reward and resource- i consumption over these rounds are close to $\bar{X}_\tau \cdot \mathbf{r}$ and $\bar{X}_\tau \cdot \mathbf{c}_i$, respectively,

with high probability. Specifically, a simple usage of Azuma-Hoeffding inequality (Lemma A.1) implies that

$$\frac{1}{\tau} \sum_{t \in [\tau]} r_t \geq \bar{X}_\tau \cdot \mathbf{r} - R_0(\tau)/\tau, \quad (4.10)$$

$$\frac{1}{\tau} \sum_{t \in [\tau]} c_{i,t} \leq \bar{X}_\tau \cdot \mathbf{c}_i + R_0(\tau)/\tau, \quad \forall i \in [d], \quad (4.11)$$

hold with probability at least $1 - \delta$, where $R_0(\tau) = O(\sqrt{\tau \log(d/\delta)})$.

Regret minimization in games. Let us apply the machinery from regret minimization in games to the repeated Lagrangian game. Consider the game matrix G of the expected game. Using Equation (4.6) and Lemma 4.2, we conclude that the minimax value of G is $v^* = \text{OPT}_{\text{LP}}(\mathbf{M}, B, T)$.

We apply Lemma 3.1, with a fixed stopping time $\tau \in [T]$. Recall that the payoff range is $b_{\max} - b_{\min} = \frac{T}{B} + 1$. Thus, with probability at least $1 - 2\delta$ it holds that

$$\lambda \in \Delta_d : \quad \bar{X}_\tau^T G \lambda \geq v^* - \frac{1}{\tau} \left(\frac{T}{B} + 1 \right) \cdot \text{reg}(T), \quad (4.12)$$

where the regret term is $\text{reg}(T) := R_{1, \delta/T}(T) + R_{2, \delta/T}(T) + 4\sqrt{2T \log(T/\delta)}$.

Crux of the proof. Let us condition on the event that (4.10), (4.11), and (4.12) hold for each $\tau \in [T]$. By the union bound, this event holds with probability at least $1 - 3\delta T$.

Let τ denote the *stopping time* of the algorithm, the first round when the total consumption of some resource exceeds its budget. Let i be the resource for which this happens; hence,

$$\sum_{t \in [\tau]} c_{i,t} > B. \quad (4.13)$$

Let us use Equation (4.12) with $\lambda = \lambda^{(i)}$, the point distribution for this resource. Then

$$\begin{aligned} \bar{X}_\tau^T G \lambda^{(i)} &= \mathcal{L}_{\mathbf{M}, B, T}(\bar{X}_\tau, \lambda^{(i)}) && \text{(by Equation (4.6))} \\ &= \bar{X}_\tau \cdot \mathbf{r} + 1 - \frac{T}{B} \bar{X}_\tau \cdot \mathbf{c}_i && \text{(by definition of Lagrange function)} \\ &\leq \frac{1}{\tau} \left(\left(\sum_{t \in [\tau]} r_t \right) - \left(\frac{T}{B} \sum_{t \in [\tau]} c_{i,t} \right) + \tau + \left(1 + \frac{T}{B} \right) R_0(\tau) \right) && \text{(plugging in (4.10) and (4.11))} \\ &\leq \frac{1}{\tau} \left(\left(\sum_{t \in [\tau]} r_t \right) + \tau - T + \left(1 + \frac{T}{B} \right) R_0(\tau) \right). && \text{(plugging in Equation (4.13))} \end{aligned}$$

Plugging this into Equation (4.12) and rearranging, we obtain

$$\sum_{t \in [\tau]} r_t \geq \tau v^* + T - \tau - \left(1 + \frac{T}{B} \right) \cdot \text{reg}(T) - \left(1 + \frac{T}{B} \right) R_0(\tau).$$

Since $v^* \leq 1$ (because $v^* = \text{OPT}_{\text{LP}}$, as we've proved above),

$$\text{REW}(\text{LagrangeBwK}) = \sum_{t \in [\tau]} r_t \geq T v^* - \left(1 + \frac{T}{B} \right) \cdot \text{reg}(T) - \left(1 + \frac{T}{B} \right) R_0(\tau).$$

The claimed regret bound (4.9) follows by Equation (4.2), completing the proof of Theorem 4.4.

ALGORITHM 2: A simple algorithm for Adversarial BwK.

input: scale parameter $\kappa > 1$, guess range $[g_{\min}, g_{\max}]$, primal and dual algorithms $\text{ALG}_1, \text{ALG}_2$
 // $\text{ALG}_1, \text{ALG}_2$ are adversarial online learning algorithms
 // with bandit feedback and full feedback, resp.
 Choose u uniformly at random from $[0, u_{\max}]$, where $u_{\max} = \log_{\kappa} \frac{g_{\max}}{g_{\min}}$.
 Guess the value of OPT_{FD} as $\widehat{g} = g_{\min} \cdot \kappa^u$.
 Run LagrangeBwK with algorithms $\text{ALG}_1, \text{ALG}_2$ and parameters $B_0 = B$ and $T_0 = \widehat{g}/(d+1)$.

5 A SIMPLE ALGORITHM FOR ADVERSARIAL BWK

We present and analyze an algorithm for Adversarial BwK which achieves $d \cdot \log T$ competitive ratio, in expectation, up to a low-order additive term. Our algorithm is very simple: we randomly guess the value of OPT_{FD} and run LagrangeBwK with parameter T_0 driven by this guess. The analysis is very different, however, since we cannot rely on the machinery from regret minimization in stochastic games. The crux of the analysis (Lemma 5.8) is re-used to analyze the high-probability algorithm in the next section.

In hindsight, the intuition for our algorithm can be explained as follows. Since LagrangeBwK builds on adversarial online learning algorithms ALG_j , it appears plausibly applicable to Adversarial BwK. We analyze it for an arbitrary parameter T_0 , and find that it performs best when T_0 is tailored to OPT_{FD} up to a constant multiplicative factor. This is precisely what our algorithm achieves using the random guess.

Our algorithm is presented as Algorithm 2. We guess the value of OPT_{FD} within a given range $[g_{\min}, g_{\max}]$. We guess OPT_{FD} uniformly on the “exponential scale”: we draw the exponent u uniformly at random, and define the guess as $\widehat{g} = g_{\min} \cdot \kappa^u$, for some scale parameter $\kappa > 1$.¹⁰ We call LagrangeBwK with $T_0 = \widehat{g}/(d+1)$. Our analysis works as long as $\text{OPT}_{\text{FD}} \leq g_{\max}$, with g_{\min} appearing in the additive term.

THEOREM 5.1. *Consider Adversarial BwK with K arms, d resources, time horizon T , and budget B . Assume that one of the arms is a null arm that has zero reward and zero resource consumption. Consider Algorithm 2 with scale parameter $\kappa > 1$. Suppose algorithms ALG_j that satisfy the regret bound (3.2) with $\delta = T^{-2}$ and regret term $R_{\delta}(T) = R_j, \delta(T)$, for any known payoff range $[b_{\min}, b_{\max}]$.*

(a) If $\text{OPT}_{\text{FD}} \leq g_{\max}$ then the expected reward of Algorithm 2 satisfies

$$\mathbb{E}[\text{REW}] \geq \frac{\text{OPT}_{\text{FD}} - g_{\min}}{(d+1) \ln \left(\frac{g_{\max}}{g_{\min}} \right)} - \text{reg} - 1, \quad (5.1)$$

where $\text{reg} = \left(1 + \frac{\text{OPT}_{\text{FD}}}{dB}\right) (R_1, \delta/T(T) + R_2, \delta/T(T))$.

(b) In particular, taking $[g_{\min}, g_{\max}] = [\sqrt{T}, T]$, we obtain

$$\mathbb{E}[\text{REW}] \geq \frac{\text{OPT}_{\text{FD}} - \sqrt{T}}{\frac{1}{2}(d+1) \ln(T)} - \text{reg} - 1. \quad (5.2)$$

Remark 5.2. One can use algorithms EXP3.P for ALG_1 and Hedge for ALG_2 , with regret bounds given by (3.4), and achieve the regret term $\text{reg} = O\left(1 + \frac{\text{OPT}_{\text{FD}}}{dB}\right) \sqrt{TK \log(Td/\delta)}$. We obtain a meaningful performance guarantee as long as, say, $\text{reg} < \text{OPT}_{\text{FD}}/2$; this requires OPT_{FD} and B to be at least $\widetilde{\Omega}(\sqrt{TK})$.

¹⁰Somewhat surprisingly, our results do not depend on the value κ . This is because the dependence on κ is captured via a normalized integral $\ln \kappa \cdot \int_0^{\log_{\kappa} x} \kappa^u du$, and this expression does not depend on κ .

Remark 5.3. We define the outcome matrices slightly differently compared to Section 4 in that we do not posit a dummy resource. Formally, we assume that the null arm has zero consumption in every resource. This is essential for case 1 (i.e., when $\tau_{\text{alg}} \leq \sigma$) in the analysis of Lemma 5.8.

Remark 5.4. The $\log(T)$ appears in the competitive ratio because the algorithm needs to guess OPT_{FD} up to a constant factor. The factor of d can be traced to a pessimistic over-estimate in (5.12).

Remark 5.5. The algorithm simplifies when $d = 1$, i.e., if there is only one resource other than the dummy resource. Then the outcome matrices have only one resource, so the dual algorithm ALG_2 is no longer needed.

Remark 5.6. The problem can be reduced to the case $d = 1$, which simplifies the algorithm, as per Remark 5.5, but increases the competitive ratio. The reduction is very simple: replacing all “true resources” (i.e., all resources other than the dummy resource) with the “maximal resource” whose consumption is the maximum over the true resources. The competitive ratio, i.e., the denominator in Equation (5.1), increases by the factor of $\frac{2d}{d+1}$. Moreover, the reduction can be wasteful if the maximal consumption (across all resources) is much larger than a “typical” consumption of each resource. The analysis compares algorithm’s reward to the benchmark for the “fake problem” with $d = 1$, then compares the said benchmark to OPT_{FD} . The former step is essentially the analysis in Section 5.1, albeit in a slightly simpler form. We omit the easy details.

If a problem instance of Adversarial BwK is actually an instance of adversarial bandits, then we recover the optimal $\tilde{O}(\sqrt{KT})$ regret. (This easily follows by examining the proof of Lemma 5.8.)

LEMMA 5.7. *Consider LagrangeBwK, with algorithms EXP3.P for ALG_1 and Hedge for ALG_2 , for an instance of Adversarial BwK with zero resource consumption. This algorithm obtains $\tilde{O}(\sqrt{KT})$ regret, for any parameters $B_0, T_0 > 0$. Accordingly, so does Algorithm 2 with any scale parameter $\kappa > 0$.*

5.1 Analysis: Proof of Theorem 5.1 and Lemma 5.7

Stopped linear program. Let us set up a linear relaxation that is suitable to the adversarial setting. The expected outcome matrix is no longer available. Instead, we use *average* outcome matrices:

$$\bar{\mathbf{M}}_\tau = \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{M}_t, \quad (5.3)$$

the average up to a given intermediate round $\tau \in [T]$. Similar to the stochastic case, the relaxation assumes that each instantaneous outcome matrix \mathbf{M}_t is equal to the average outcome matrix $\bar{\mathbf{M}}_\tau$. What is different now is that the relaxation depends on τ : using $\bar{\mathbf{M}}_\tau$ is tantamount to stopping precisely at this round.

With this intuition in mind, for a particular end-time τ we consider the linear program (4.1), parameterized by the time horizon τ and the average outcome matrix $\bar{\mathbf{M}}_\tau$. Its value, $\text{OPT}_{\text{LP}}(\bar{\mathbf{M}}_\tau, B, \tau)$, represents the per-round expected reward, so it needs to be scaled by the factor of τ to obtain the total expected reward. Finally, we maximize over τ . Thus, our linear relaxation for Adversarial BwK is defined as follows:

$$\text{OPT}_{\text{LP}}^{[T]} := \max_{\tau \in [T]} \tau \cdot \text{OPT}_{\text{LP}}(\bar{\mathbf{M}}_\tau, B, \tau) \geq \text{OPT}_{\text{FD}}. \quad (5.4)$$

The inequality in (5.4) is proved in the appendix (Section A.3).

Regret bounds for ALG_j . Since each algorithm ALG_j , $j \in \{1, 2\}$ satisfies regret bound (3.2) with $\delta = T^{-2}$ and $R_\delta(T) = R_{j,\delta}(T)$, it also satisfies a stronger version (3.3) with the same parameters. Recall from (4.7) that the payoff range is $[b_{\min}, b_{\max}] = [-\frac{T_0}{B} + 1, 2]$. For succinctness, let $U_j(T|T_0) = (1 + \frac{T_0}{B})R_{j,\delta/T}(T)$ denote the respective regret term in (3.3).

Let us apply these regret bounds to our setting. Let $a_t \in [K]$ and $i_t \in [d]$ be, respectively, the chosen arm and resource in round t . We represent the outcomes as vectors over arms: $\mathbf{r}_t, \mathbf{c}_{t,i} \in [0, 1]^K$ denote, resp, reward vector and resource- i consumption vector for a given round t . Recall that the round- t payoffs in LagrangeBwK are given by the Lagrange function $\mathcal{L}_t := \mathcal{L}_{M_t, B, T_0}$ such that

$$\mathcal{L}_t(a, i) = r_t(a) + 1 - \frac{T_0}{B} c_{t,i}(a) \quad (5.5)$$

for each arm a and resource i . Consider the total Lagrangian payoff at a given round $\tau \in [T]$:

$$\sum_{t \in [\tau]} \mathcal{L}_t(a_t, i_t) = \text{REW}_\tau + \tau - W_\tau, \quad (5.6)$$

where $\text{REW}_\tau = \sum_{t \in [\tau]} r_t(a_t)$ is the total reward up to round τ , and $W_\tau = \frac{T_0}{B} \sum_{t \in [\tau]} c_{t,i_t}(a_t)$ is the *consumption term*. The regret bounds sandwich (5.6) from above and below:

$$\left(\max_{a \in [K]} \sum_{t \in [\tau]} \mathcal{L}_t(a, i_t) \right) - U_1(T|T_0) \leq \text{REW}_\tau + \tau - W_\tau \leq \left(\min_{i \in [d]} \sum_{t \in [\tau]} \mathcal{L}_t(a_t, i) \right) + U_2(T|T_0). \quad (5.7)$$

This holds for all $\tau \in [T]$, with probability at least $1 - 2\delta$. The first inequality in (5.7) is due to the primal algorithm, and the second is due to the dual algorithm. Call them *primal* and *dual* inequality, respectively.

Crux of the proof. We condition on the event that (5.7) holds for all $\tau \in [T]$, which we call the *clean event*. The crux of the analysis is encapsulated in the following lemma, which analyzes an execution of LagrangeBwK with an arbitrary parameter T_0 under the clean event.

LEMMA 5.8. *Consider an execution of LagrangeBwK with $B_0 = B$ and an arbitrary parameter T_0 such that the clean event holds. Fix an arbitrary round $\sigma \in [T]$, and consider the LP value relative to this round:*

$$f(\sigma) := \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_\sigma, B, \sigma). \quad (5.8)$$

The algorithm's reward up to round σ satisfies

$$\text{REW}_\sigma \geq \min(T_0, \sigma \cdot f(\sigma) - dT_0) - (U_1(T|T_0) + U_2(T|T_0)). \quad (5.9)$$

Taking σ to be the maximizer in (5.4), algorithm's reward satisfies

$$\text{REW} \geq \min(T_0, \text{OPT}_{\text{FD}} - dT_0) - (U_1(T|T_0) + U_2(T|T_0)). \quad (5.10)$$

Equation (5.9) is used, with a different σ , for the high-probability analysis in Section 6.

PROOF. Let τ_{alg} be the stopping time of the algorithm. We consider two cases, depending on whether some resource is exhausted at time σ . In both cases, we focus on the round $\min(\tau_{\text{alg}}, \sigma)$.

Case 1: $\tau_{\text{alg}} \leq \sigma$ and some resource is exhausted. Let us focus on round $\tau = \tau_{\text{alg}}$. If i is the exhausted resource, then $\sum_{t \in [\tau]} c_{t,i}(a_t) > B$. Let us apply the dual inequality in (5.7) for this resource:

$$\begin{aligned} \text{REW}_\tau + \tau - W_\tau - U_2(T|T_0) &\leq \sum_{t \in [\tau]} \mathcal{L}_t(a_t, i) \\ &= \text{REW}_\tau + \tau - \frac{T_0}{B} \sum_{t \in [\tau]} c_{t,i}(a_t) \\ &\leq \text{REW}_\tau + \tau - T_0. \end{aligned}$$

It follows that $W_\tau \geq T_0 - U_2(T|T_0)$.

Now, let us apply the primal inequality in (5.7) for the null arm. Recall that the reward and consumption for this arm is 0, so $\mathcal{L}_t(\text{null}, i_t) = 1$ for each round t . Therefore,

$$\text{REW}_\tau + \tau - W_\tau + U_1(T|T_0) \geq \sum_{t \in [\tau]} \mathcal{L}_t(\text{null}, i_t) = \tau.$$

We conclude that $\text{REW}_\tau \geq W_\tau - U_1(T|T_0) \geq T_0 - U_1(T|T_0) - U_2(T|T_0)$.

Case 2: $\tau_{\text{alg}} \geq \sigma$. Let us focus on round σ . Consider the linear program $\text{LP}_{\overline{M}_\sigma, B, \sigma}$, and let $X^* \in \Delta_K$ be an optimal solution to this LP. The primal inequality in (5.7) implies that

$$\begin{aligned} \text{REW}_\sigma + \sigma - W_\sigma + U_1(\sigma) &\geq \max_{a \in [K]} \sum_{t \in [\sigma]} \mathcal{L}_t(a, i_t) \\ &\geq \sum_{t \in [\sigma]} \sum_{a \in [K]} X^*(a) \mathcal{L}_t(a, i_t) \\ &= \sigma + \sum_{t \in [\sigma]} X^* \cdot \mathbf{r}_t - \frac{T_0}{B} \sum_{t \in [\sigma]} X^* \cdot \mathbf{c}_{t, i_t} \\ \text{REW}_\sigma &\geq \sigma \cdot f(\sigma) - \frac{T_0}{B} \sum_{t \in [\sigma]} X^* \cdot \mathbf{c}_{t, i_t} - U_1(T|T_0). \end{aligned} \quad (5.11)$$

In the last inequality we used the fact that $\sum_{t \in [\sigma]} X^* \cdot \mathbf{r}_t = \sigma \cdot f(\sigma)$ by optimality of X^* .

$\sum_{t \in [\sigma]} X^* \cdot \mathbf{c}_{t, i} \leq B$ for each resource i , since X^* is a feasible solution for $\text{OPT}_{\text{LP}}(\overline{M}_\sigma, B, \sigma)$. Then,

$$\sum_{t \in [\sigma]} X^* \cdot \mathbf{c}_{t, i_t} \leq \sum_{i \in [d]} \sum_{t \in [\sigma]} X^* \cdot \mathbf{c}_{t, i} \leq dB. \quad (5.12)$$

Plugging (5.12) into (5.11), we conclude that $\text{REW}_\sigma \geq \sigma \cdot f(\sigma) - dT_0 - U_1(T|T_0)$.

Conclusions from the two cases imply (5.10), as claimed. \square

Wrapping up (the easy version). $\text{OPT}_{\text{FD}} \in [g_{\min}, g_{\max}]$, then some guess \widehat{g} is approximately correct:

$$\text{OPT}_{\text{FD}}/\kappa \leq \widehat{g} \leq \text{OPT}_{\text{FD}}. \quad (5.13)$$

By Lemma 5.8, the algorithm's execution with this guess, assuming the clean event, satisfies (5.10), where, recalling that $T_0 = \widehat{g}/(d+1)$, we have

$$\min(T_0, \text{OPT}_{\text{FD}} - dT_0) \geq \frac{\text{OPT}_{\text{FD}}}{\kappa(d+1)} \quad \text{and} \quad T_0 \leq \frac{\text{OPT}_{\text{FD}}}{d+1}.$$

The regret term for this guess is

$$\text{reg} = U_1(T|T_0) + U_2(T|T_0) \leq \left(1 + \frac{\text{OPT}_{\text{FD}}}{(d+1)B}\right) (R_1, \delta_{/T}(T) + R_2, \delta_{/T}(T)).$$

To complete the proof of (5.1) (with a much larger constant in the denominator), note that we obtain a suitable guess \widehat{g} with probability $1/\lceil \log_\kappa \frac{g_{\max}}{g_{\min}} \rceil$.

Wrapping up (optimizing the constants). Let us go beyond the ‘‘approximately right guess’’ in (5.13), and account for contributions of *every* guess. In other words, let us integrate over the guesses.

Assume that OPT_{FD} lies in the guess range $[g_{\min}, g_{\max}]$. Recall that the algorithm samples u uniformly at random from the interval $[0, u_{\max}]$. Write

$$\begin{aligned} T_0 &= T_0(u) = g_{\min} \cdot \kappa^u / (d + 1), \\ \text{reg} &= \text{reg}(u) = U_1(T|T_0(u)) + U_2(T|T_0(u)), \\ \Lambda(u) &= \min(T_0(u), \max(0, \text{OPT}_{\text{FD}} - dT_0(u))). \end{aligned}$$

Then by Lemma 5.8, for a particular choice of u , the algorithm's reward satisfies

$$\mathbb{E}[\text{REW} \mid u] \geq \Lambda(u) - \text{reg}(u) - 1, \quad (5.14)$$

where the '-1' term accounts for the complement of the "clean event".

The $\Lambda(u)$ term can be split into three cases as follows:

$$\Lambda(u) = \begin{cases} T_0(u) & \text{if } 0 \leq u \leq \log_{\kappa} \frac{\text{OPT}_{\text{FD}}}{g_{\min}}, \\ \text{OPT}_{\text{FD}} - dT_0(u) & \text{if } \log_{\kappa} \frac{\text{OPT}_{\text{FD}}}{g_{\min}} < u \leq \log_{\kappa} \left(\frac{d+1}{d} \cdot \frac{\text{OPT}_{\text{FD}}}{g_{\min}} \right), \\ 0 & \text{otherwise.} \end{cases} \quad (5.15)$$

So, we are only interested in $u \leq u^* := \log_{\kappa} \left(\frac{\text{OPT}_{\text{FD}}}{g_{\min}} \right)$.

Integrating the right-hand side of (5.14) over u , we obtain:

$$\mathbb{E}[\text{REW}] \geq \frac{1}{u_{\max}} \int_0^{u^*} \mathbb{E}[\text{REW} \mid u] du = \frac{1}{u_{\max}} \int_0^{u^*} (\Lambda(u) - \text{reg}(u) - 1) du, \quad (5.16)$$

where $u_{\max} = \log_{\kappa} \frac{g_{\max}}{g_{\min}}$ as per the algorithm's specification.

Using (5.15) and omitting the easy details, the main term $\Lambda(u)$ integrates as follows:

$$\int_0^{u^*} \Lambda(u) du \geq \frac{\text{OPT}_{\text{FD}} - g_{\min}}{(\ln \kappa)(d + 1)}. \quad (5.17)$$

(We've only used the first "regime" in (5.15). As for the second "regime" in (5.15), integrating over it only improves (5.17) by a small additive term.)

To handle the regret term $\text{reg}(u)$, note that it is non-decreasing with u , so

$$\frac{1}{u_{\max}} \int_0^{u^*} \text{reg}(u) du \leq \frac{u^* \cdot \text{reg}(u^*)}{u_{\max}} \leq \text{reg}(u^*).$$

Plugging this into (5.16), we obtain

$$\mathbb{E}[\text{REW}] \geq \frac{1}{u_{\max}} \left(\frac{\text{OPT}_{\text{FD}} - g_{\min}}{(\ln \kappa)(d + 1)} \right) - \text{reg}(u^*) - 1. \quad (5.18)$$

Recalling that $T_0 = T_0(u^*) = \text{OPT}_{\text{FD}} / (d + 1)$, we have

$$\text{reg}(u^*) = U_1(T|T_0) + U_2(T|T_0) \leq \left(1 + \frac{\text{OPT}_{\text{FD}}}{(d + 1)B} \right) (R_1, \delta/T(T) + R_2, \delta/T(T)).$$

Finally, because of the $-g_{\min}$ term in (5.18), the assumption $\text{OPT}_{\text{FD}} \geq g_{\min}$ is redundant. This completes the proof of Theorem 5.1(a).

Proof Sketch of Lemma 5.7. Recall that in the adversarial bandit setting we have $c_{i,t} = 0$ for every $i \in [d]$ and every $t \in [T]$. We re-analyze Lemma 5.8 with $\sigma = T$. Notice that case 1 never occurs. Thus, we obtain Equation (5.11) in case 2. Note that $\frac{T_0}{B} \sum_{t \in [\sigma]} \mathbf{X}^* \cdot \mathbf{c}_{t,i_t} = 0$ since $c_{i,t} = 0$. Therefore, we obtain

$$\text{REW}_T \geq T \cdot f(T) - U_1(T|T_0).$$

We now argue that $T \cdot f(T) = \max_{a \in [K]} \sum_{t \in [T]} r_t(a)$. Let X^* be the optimal distribution over the arms. Thus, $\sum_{t \in [T]} X^* \cdot \mathbf{r}_t = T \cdot f(T)$. Note that since $c_{i,t} = 0$ the only constraint on X^* is that it lies in Δ_K . Therefore, the maximizer is a point distribution on $\max_{a \in [K]} \sum_{t \in [T]} r_t(a)$. This proof does not rely on any specific value for B_0, T_0 . The payoff range is $[b_{\max}, b_{\min}] = [1, 2]$, so $U_1(T|T_0) = \tilde{O}(\sqrt{KT})$.

6 HIGH-PROBABILITY ALGORITHM FOR ADVERSARIAL BWK

We recover the $O(d \log T)$ competitive ratio for Adversarial BwK, but with high probability rather than merely in expectation. Our algorithm uses LagrangeBwK as a subroutine, and re-uses the adversarial analysis thereof (Lemma 5.8). We do not optimize the regret term or the constant in the competitive ratio.

The algorithm is considerably more complicated compared to Algorithm 2. Instead of making one random guess \hat{g} for the value of $\text{OPT}_{\text{LP}}^{[T]}$, we iteratively refine this guess over time. The algorithm proceeds in phases. In the beginning of each phase, we start a fresh instance of LagrangeBwK with parameter T_0 defined by the current value of \hat{g} .¹¹ We update the guess \hat{g} in each round (in a way specified later), and stop the phase once \hat{g} becomes too large compared to its initial value in this phase. We invoke LagrangeBwK with a rescaled budget $B_0 = B/\Theta(\log T)$. Within each phase, we simulate the BwK problem with budget B_0 : we stop LagrangeBwK once the consumption of some resource in this phase exceeds B_0 . For the remainder of the phase, we play the null arm with probability $1 - \gamma_0$ and do uniform exploration with the remaining probability, for some parameter $\gamma_0 \in (0, 1)$ (here and elsewhere, *uniform exploration* refers to choosing each action with equal probability). The pseudocode is summarized in Algorithm 3.

ALGORITHM 3: High-probability algorithm for Adversarial BwK.

```

input: scale parameter  $\kappa$ , exploration parameter  $\gamma_0$ , primal algorithm  $\text{ALG}_1$ , dual algorithm  $\text{ALG}_2$ 
//  $\text{ALG}_1, \text{ALG}_2$  are adversarial online learning algorithms
// with bandit feedback and full feedback, resp.
Initialize  $\hat{g} = 1$ .
for each phase do
  Start a fresh instance ALG of LagrangeBwK
  with parameters  $B_0 = B/2\lceil \log_\kappa T \rceil$  and  $T_0 = \hat{g}/(3d\lceil \log_\kappa T \rceil)$ .
  Define  $\hat{g}_{\text{old}} := \hat{g}$ .
  for each round in this phase do
    Recompute the guess  $\hat{g}$ 
    if  $\hat{g} > \kappa \cdot \hat{g}_{\text{old}}$  then start a new phase
    if consumption of all resources in this phase does not exceed  $B_0$  then
      | Play the action chosen by ALG, observe the outcome and report it back to ALG.
    else
      | Choose the null arm with probability  $1 - \gamma_0$ , do uniform exploration otherwise

```

To complete the algorithm's specification, let us define how to update the guess \hat{g} in each round t . The guess, denoted \hat{g}_t , is an estimate for $\text{OPT}_{\text{LP}}^{[t]}$, -as defined in (5.4). We form this estimate using a standard *inverse propensity scoring (IPS)* technique. Let \mathbf{p}_t and a_t be, resp., the distribution and the arm chosen by the primal algorithm in round t . The instantaneous outcome matrix \mathbf{M}_t is

¹¹The idea of restarting the algorithm in each phase is similar to the standard "doubling trick" in the online machine learning literature, but much more delicate in our setting.

estimated by matrix $\mathbf{M}_t^{\text{ips}} \in [0, \infty)^{K \times d}$ such that each row $\mathbf{M}_t^{\text{ips}}(a)$ is defined as follows:

$$\mathbf{M}_t^{\text{ips}}(a) := \mathbf{1}_{\{a_t=a\}} \frac{1}{p_t(a_t)} \mathbf{M}_t(a). \quad (6.1)$$

For a given end-time τ , the average outcome matrix $\overline{\mathbf{M}}_\tau$ from (5.3) is estimated as

$$\overline{\mathbf{M}}_\tau^{\text{ips}} := \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{M}_t^{\text{ips}}. \quad (6.2)$$

Finally, we plug this estimate into (5.3) and define

$$\widehat{g}_t := \max_{\tau \in [t]} \tau \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_\tau^{\text{ips}}, B, \tau). \quad (6.3)$$

For the analysis, we will assume that the primal algorithm does some uniform exploration:

$$p_t(a) \geq \gamma > 0 \quad \text{for each arm } a \in [K] \text{ and each round } t \in [T]. \quad (6.4)$$

THEOREM 6.1. *Consider Adversarial BwK with K arms, d resources, time horizon T , and budget B . Let $\delta > 0$ be the failure probability parameter. Assume that $B > 5K T^{3/4} \log(2T^2)$. Suppose that one of the arms is a null arm that has zero reward and zero resource consumption.*

Consider Algorithm 3 with parameters $\kappa = 2$ and $\gamma_0 = T^{-1/4}$. Assume that each algorithm ALG_j , $j \in \{1, 2\}$, satisfies the regret bound (3.2) with payoff range $[b_{\min}, b_{\max}] = [-\frac{T}{B} + 1, 2]$ and regret term $R_\delta(T) = R_{j,\delta}(T)$. Assume that the primal algorithm ALG_1 satisfies (6.4) with parameter $\gamma \geq T^{-1/4}$.

Then the total reward REW collected by Algorithm 3 satisfies

$$\Pr \left[\text{REW} \geq \frac{\text{OPT}_{\text{FD}}}{O(d \log T)} - O(\text{reg}) \right] \geq 1 - O(\delta T), \quad (6.5)$$

where the regret term is $\text{reg} = \frac{T}{B} (K T^{3/4} \log^{1/2}(\frac{1}{\delta}) + R_{1,\delta/T}(T) + R_{2,\delta/T}(T))$.

Remark 6.2. Using algorithms EXP3.P for ALG_1 and Hedge for ALG_2 , we can achieve (6.5) with

$$\text{reg} = O \left(\frac{TK}{B} \right) T^{3/4} \sqrt{\log(T/\delta)}.$$

This is because EXP3.P, with appropriately modified uniform exploration term $\gamma = T^{-1/4}$, satisfies the regret bound (3.2) with $R_\delta(\tau) = O(T^{3/4}) \sqrt{K \log \frac{T}{\delta}}$, and for Hedge we can (still) use Equation (3.4). The theorem is meaningful whenever, say, $\text{reg} < \text{OPT}_{\text{FD}}/2$. The latter requires $\text{OPT}_{\text{FD}} \cdot \frac{B}{K} > \widetilde{\Omega}(T^{7/4})$.

Remark 6.3. Like in Theorem 5.1, we posit that the null arm does not consume any resources.

Remark 6.4. For the sake of intuition, let us clarify the choice of parameters B_0 and T_0 in the algorithm. First, $\lceil \log_\kappa T \rceil$ appears in both B_0 and T_0 , because it is an upper bound on the number of phases. Second, d is needed in T_0 to counteract the dependence on d in Lemma 5.8, the adversarial analysis of LagrangeBwK. Third, the $\frac{1}{2}$ appears in B_0 because we allow half of the budget to be spent on uniform exploration. Finally, the $\frac{1}{3}$ in T_0 is needed to enable a specific step deep down in the analysis, the transition from Equation (6.17) to Equation (6.19).

PROOF SKETCH OF THEOREM 6.1. The proof consists of several steps. First, we argue that the guess \widehat{g}_t is close to $\text{OPT}_{\text{LP}}^{[t]}$ with high probability. This argument only relies on the uniform exploration property (6.3) and the definition of IPS estimators, not on any properties of the algorithm. We immediately obtain concentration for the average outcome matrices. With some work, we derive concentration on the respective LP-values.

Next, we focus on a particular phase in the execution of the algorithm. We say that a phase is *full* if the stopping condition $\widehat{g}_t > \kappa \cdot \widehat{g}_{\text{old}}$ has fired. We focus on the last full phase. We prove there

is enough reward to be collected in this phase. Essentially, letting τ_1, τ_2 be, resp., the start and end time of this phase, we consider the BwK problem restricted to time interval $[\tau_1, \tau_2]$, and lower-bound the LP-value of this problem in terms of the LP-value of the original problem. Finally, we use the adversarial analysis of LagrangeBwK (Lemma 5.8) to guarantee that our algorithm actually collects that value.

Because of the stopping condition $\widehat{g}_t > \kappa \cdot \widehat{g}_{\text{old}}$, there can be at most $\lceil \log_\kappa T \rceil$ phases. Therefore, rescaling the budget to $B_0/2 \lceil \log_\kappa T \rceil$ guarantees that the algorithm consumes at most $B/2$ of the budget. We then argue that, with high-probability, the additional uniform exploration in each phase, consumes a budget of at most $B/2$ with high-probability. Thus, the algorithm never runs out of budget. \square

6.1 Full Proof of Theorem 6.1

We now describe the full proof of Theorem 6.1, following the plan outlined in the proof sketch. We decompose the analysis into several distinct pieces, present them one by one, and then show how to put them together. Each piece is presented as a lemma, with appropriate notation and intuition.

For clarity, most of the analysis is presented for an arbitrary parameter $\kappa > 1$ in the algorithm, as long as it is an absolute constant, and an arbitrary parameter γ in Equation (6.4). We only plug in $\kappa = 2$ and $\gamma \geq T^{-1/4}$ in the very end, in Equation (6.19) and right after.

Extended notation. To argue about a given phase, we extend some of our notation to refer to arbitrary time intervals, not just $[1, \tau]$. In what follows, fix time interval $[\tau_1, \tau_2]$, and let $\Delta\tau = \tau_2 - \tau_1 + 1$. Let

$$\overline{\mathbf{M}}_{[\tau_1, \tau_2]} := \frac{1}{\Delta\tau} \sum_{t=\tau_1}^{\tau_2} \mathbf{M}_t,$$

$$\overline{\mathbf{M}}_{[\tau_1, \tau_2]}^{\text{ips}} := \frac{1}{\Delta\tau} \sum_{t=\tau_1}^{\tau_2} \mathbf{M}_t^{\text{ips}}$$

be, resp., the average outcome matrix and its IPS-estimate on this time interval. Define

$$\text{OBJ}([\tau_1, \tau_2]) := \Delta\tau \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{[\tau_1, \tau_2]}, B, \Delta\tau), \quad (6.6)$$

$$\text{OBJ}^{\text{ips}}([\tau_1, \tau_2]) := \Delta\tau \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{[\tau_1, \tau_2]}^{\text{ips}}, B, \Delta\tau). \quad (6.7)$$

We use short-hand $\text{OBJ}(\tau_2) = \text{OBJ}([1, \tau_2])$ and $\text{OBJ}^{\text{ips}}(\tau_2) = \text{OBJ}^{\text{ips}}([1, \tau_2])$. We think of these quantities, resp., as the LP-objective given the stopping time at τ_2 , and the IPS-estimate thereof. Recall that

$$\text{OPT}_{\text{LP}}^{[\tau]} := \max_{t \in [\tau]} \text{OBJ}(t). \quad (6.8)$$

$$\widehat{g}_\tau := \max_{t \in [\tau]} \text{OBJ}^{\text{ips}}(t). \quad (6.9)$$

Uniform exploration does not exhaust budget. The uniform exploration in Algorithm 3 happens for at most $\gamma_0 T$ rounds in expectation, and therefore for at most $\gamma_0 T + 3\sqrt{\gamma_0 T \ln(1/\delta)}$ rounds with probability at least $1 - \delta$.¹² It does not consume more than $B/2$ units of each resource, since $\gamma_0 = T^{-1/4}$ and $B > 4T^{3/4}$.

IPS estimators are good. We argue that, essentially, the guess \widehat{g}_τ is close to $\text{OPT}_{\text{LP}}^{[\tau]}$ with high probability. To this end, we prove that $\text{OBJ}(\tau)$ is close to its IPS estimator, for any given $\tau \in [T]$.

¹²By an easy application of Chernoff-Hoeffding bounds (Lemma A.2).

LEMMA 6.5. *With probability at least $1 - d\delta T$ it holds that*

$$\forall \tau \in [T] \quad |\text{OBJ}^{\text{ips}}(\tau) - \text{OBJ}(\tau)| \leq \text{DEV}(\tau) := \left(1 + \frac{2\text{OBJ}(\tau)}{B}\right) \frac{K}{\gamma} \sqrt{2\tau \log \frac{T}{\delta}}, \quad (6.10)$$

If the event (6.10) holds, then \widehat{g}_τ and $\text{OPT}_{\text{LP}}^{[\tau]}$ are indeed close:

$$\forall \tau \in [T] \quad \left| \widehat{g}_\tau - \text{OPT}_{\text{LP}}^{[\tau]} \right| \leq \text{DEV}_{\max} := \frac{KT}{\gamma B} \sqrt{18T \log \frac{T}{\delta}} \quad (6.11)$$

The proof of this lemma is deferred Section 6.3. It only relies on the uniform exploration property (6.3) and the definition of IPS estimators, not on anything that the algorithm does. A somewhat subtle point is to derive concentration on the respective LP-values from concentration of the average outcome matrices.

IPS estimates do not change too fast. We use the phase-stopping condition in the algorithm to argue that algorithm's guesses \widehat{g}_t and estimates $\text{OBJ}^{\text{ips}}(t)$ do not change too fast.

LEMMA 6.6. *Consider a full phase in the execution of the algorithm. Let τ be the first round in this phase, let τ' be any other round in this phase, and let τ'' be any round in the next phase. Then*

$$\widehat{g}_{\tau'} \leq \kappa \cdot \widehat{g}_\tau < \widehat{g}_{\tau''}.$$

PROOF. The first inequality holds because the phase-stopping condition did not fire at round τ' . For the second inequality, let t denote the first round in the next phase. Then

$$\begin{aligned} \widehat{g}_{\tau''} &\geq \widehat{g}_t && \text{(since } (\widehat{g}_t) \text{ is monotone by Equation (6.9))} \\ &> \kappa \cdot \widehat{g}_\tau && \text{(by the phase-stopping condition).} \end{aligned} \quad \square$$

CLAIM 6.7. *$\widehat{g}_t = \text{OBJ}^{\text{ips}}(t) > \text{OBJ}^{\text{ips}}(t-1)$ for any round t such that $\widehat{g}_t > \widehat{g}_{t-1}$. The latter condition holds, in particular, if t is the first round in some phase.*

PROOF. The claim follows from Equation (6.9) and the phase-stopping condition. In particular, if t is the first round in some phase and $\widehat{g}_t = \widehat{g}_{t-1}$, then this phase would have started earlier. \square

CLAIM 6.8. *For any round t , we have $\widehat{g}_{t+1} - \widehat{g}_t \leq \text{OBJ}^{\text{ips}}(t+1) - \text{OBJ}^{\text{ips}}(t) \leq K/\gamma$.*

PROOF. Fix round t . If $\widehat{g}_{t+1} > \widehat{g}_t$, then $\widehat{g}_{t+1} = \text{OBJ}^{\text{ips}}(t+1)$ by Claim 6.7. Since $\widehat{g}_t \geq \text{OBJ}^{\text{ips}}(t)$ by Equation (6.9), it follows that

$$\widehat{g}_{t+1} - \widehat{g}_t \leq \text{OBJ}^{\text{ips}}(t+1) - \text{OBJ}^{\text{ips}}(t). \quad (6.12)$$

Let us analyze the right-hand side in Equation (6.12). Recall the IPS-estimate matrices defined in Equations (6.1) and (6.2): the round- t matrix $\mathbf{M}_t^{\text{ips}}$ and the time-average matrix $\overline{\mathbf{M}}_t^{\text{ips}}$. Let \mathbf{X}^* denote the optimal solution to the LP induced by $\overline{\mathbf{M}}_{t+1}^{\text{ips}}$. This is the LP that determines $\text{OBJ}^{\text{ips}}(t+1)$, as per Equation (6.7). So, $\text{OBJ}^{\text{ips}}(t+1) = \sum_{\tau=1}^{t+1} \mathbf{X}^* \cdot \mathbf{r}_\tau^{\text{ips}}$, where $\mathbf{r}_\tau^{\text{ips}}$ is the reward vector in $\mathbf{M}_\tau^{\text{ips}}$.

Recall that the constraint in this LP is $\mathbf{X} \cdot \sum_{\tau=1}^{t+1} \mathbf{c}_\tau^{\text{ips}} \leq \mathbf{B}$. Since \mathbf{X}^* satisfies this constraint and $\mathbf{c}_{t+1}^{\text{ips}} \geq \mathbf{0}$, we have $\mathbf{X}^* \cdot \sum_{\tau=1}^t \mathbf{c}_\tau^{\text{ips}} \leq \mathbf{B}$. So, \mathbf{X}^* is also feasible to the LP induced by $\overline{\mathbf{M}}_t^{\text{ips}}$. It follows that $\text{OBJ}^{\text{ips}}(t) \geq \sum_{\tau=1}^t \mathbf{X}^* \cdot \mathbf{r}_\tau^{\text{ips}}$.

Putting this together, the right-hand side of Equation (6.12) is at most $\mathbf{X}^* \cdot \mathbf{r}_{t+1}^{\text{ips}}$. This is at most K/γ , since the IPS-estimated reward of each arm is at most $1/\gamma$ by Equation (6.4). \square

Last full phase offers sufficient rewards. Recall that a phase in the execution of the algorithm is called *full* if the stopping condition $\widehat{g}_t > \kappa \cdot \widehat{g}_{\text{old}}$ has fired. We focus on the last full phase; let $\tau_{\text{start}}, \tau_{\text{end}}$ denote the first and last time-steps of this phase. We prove there is enough reward to be collected in this phase.

Let τ^* denote the maximizer in Equation (5.4) which we interpret as the *optimal stopping time*. Essentially, we compare the LP value for the time interval $[\tau_{\text{start}}, \tau_{\text{end}}]$ with the LP value for the time interval $[1, \tau^*]$. The former is expressed as $\text{OBJ}([\tau_{\text{start}}, \tau_{\text{end}}])$ and the latter as $\text{OPT}_{\text{LP}}^{[T]}$. Note that the time horizon T lies in the subsequent phase (so we can apply Lemma 6.6).

LEMMA 6.9. *Consider a run of the algorithm such that event (6.10) holds. Then*

$$\text{OBJ}([\tau_{\text{start}}, \tau_{\text{end}}]) \geq \left(\frac{1}{\kappa} - \frac{1}{\kappa^2} \right) \text{OPT}_{\text{LP}}^{[T]} - O(\text{DEV}_{\text{max}}). \quad (6.13)$$

The proof of this lemma is deferred to Section 6.2.

Adversarial analysis of LagrangeBwK. Let us plug in the adversarial analysis of LagrangeBwK, as encapsulated in Lemma 5.8. We focus on the last full phase in the execution. We interpret it as an execution of algorithm LagrangeBwK with parameters B_0, T_0 on an instance of Adversarial BwK with budget B_0 that starts at round τ_{start} of the original problem. Let $\widehat{g} = \widehat{g}_{\tau_{\text{start}}}$ be the guess at the first round of the phase. Then the parameters are $B_0 = B/\text{ratio}$ and $T_0 = \widehat{g}/(3d \cdot \text{ratio})$, where $\text{ratio} = \lceil \log_{\kappa} T \rceil$.

We apply Lemma 5.8 for round $\sigma = \tau_{\text{end}} - \tau_{\text{start}} + 1$ in the execution of LagrangeBwK. Restated in our notation, $f(\sigma)$ in Lemma 5.8 becomes

$$f(\sigma) = \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{[\tau_{\text{start}}, \tau_{\text{end}}]}, B_0, \sigma).$$

Thus, we obtain that with probability at least $1 - \delta$ we have

$$\text{REW} \geq \sum_{t=\tau_{\text{start}}}^{\tau_{\text{end}}} r_t(a_t) \geq \min \left(\frac{\widehat{g}}{3d \lceil \log_{\kappa} T \rceil}, \sigma f(\sigma) - \frac{d\widehat{g}}{3d \lceil \log_{\kappa} T \rceil} \right) - \text{reg}(T), \quad (6.14)$$

where the regret term is $\text{reg}(T) := (1 + \frac{T}{B}) (R_1, \delta/T(T) + R_2, \delta/T(T))$.

Rescaling the budget. Since we use rescaled budget B_0 , we need to connect the corresponding LP-values to those for the original budget B . We use the following general fact, observed in Agrawal and Devanur [5]: for any outcome matrix \mathbf{M} , budget B , time horizon T , and rescaling factor $\psi \in (0, 1]$,

$$\text{OPT}_{\text{LP}}(\mathbf{M}, \psi B, T) \geq \psi \cdot \text{OPT}_{\text{LP}}(\mathbf{M}, B, T). \quad (6.15)$$

This holds because an optimal solution $\boldsymbol{\mu}$ to $\text{LP}_{\mathbf{M}, B, T}$, the vector $\psi \boldsymbol{\mu}$ is feasible to $\text{LP}_{\mathbf{M}, \psi B, T}$.

Putting it all together. Let us show how to complete the proof of Theorem 6.1 using the tools derived above. Throughout, we condition on the high-probability events in Lemma 6.5 and Equation (6.14).

Recall that $[\tau_{\text{start}}, \tau_{\text{end}}]$ denotes the last full phase, and let \widehat{g} denote the guess at the beginning of this phase. Recall that $\widehat{g} = \text{OBJ}^{\text{IPB}}(\tau_{\text{start}})$.

From Equation (6.15) we have that $\sigma f(\sigma) \geq \frac{1}{\text{ratio}} \text{OBJ}([\tau_{\text{start}}, \tau_{\text{end}}])$ since $B_0 = \frac{B}{\text{ratio}}$. Combining this with Equation (6.14) we obtain

$$\text{REW} = \sum_{t=\tau_{\text{start}}}^{\tau_{\text{end}}} r_t(a_t) \geq \frac{1}{\text{ratio}} \min \left(\frac{\widehat{g}}{3d}, \text{OBJ}([\tau_{\text{start}}, \tau_{\text{end}}]) - \frac{d\widehat{g}}{3d} \right) - \text{reg}(T). \quad (6.16)$$

By Lemma 6.9, we can re-write Equation (6.16) as

$$\text{REW} \geq \frac{1}{\text{ratio}} \min \left(\frac{\widehat{g}}{3d}, \left(\frac{1}{\kappa} - \frac{1}{\kappa^2} \right) \text{OPT}_{\text{LP}}^{[T]} - \frac{\widehat{g}}{3} \right) - \text{reg}(T) - O(\text{DEV}_{\max}). \quad (6.17)$$

Let us characterize how the guess \widehat{g} deviates from $\text{OPT}_{\text{LP}}^{[T]}$:

$$\text{OPT}_{\text{LP}}^{[T]} / \kappa^2 - O(\text{DEV}_{\max}) \leq \widehat{g} \leq \text{OPT}_{\text{LP}}^{[T]} / \kappa + O(\text{DEV}_{\max}). \quad (6.18)$$

To prove the upper bound in Equation (6.18),

$$\begin{aligned} \widehat{g} &\leq \widehat{g}_T / \kappa && \text{(by Lemma 6.6)} \\ &\leq \text{OPT}_{\text{LP}}^{[T]} / \kappa + \text{DEV}_{\max} / \kappa && \text{(by Lemma 6.5).} \end{aligned}$$

For the lower bound in Equation (6.18), we observe that

$$\begin{aligned} \widehat{g}_T &\leq \kappa \cdot \widehat{g}_{\tau_{\text{end}}+1} && \text{(a phase didn't start at } T) \\ &\leq \kappa (\widehat{g}_{\tau_{\text{end}}} + K/\gamma) && \text{(By Claim 6.8)} \\ &\leq \kappa (\kappa \cdot \widehat{g} + K/\gamma) && \text{(a phase didn't start at } \tau_{\text{end}}) \\ \widehat{g} &\geq \widehat{g}_T / \kappa^2 - \text{DEV}_{\max} / \kappa^2 - K/\gamma \\ &\geq \text{OPT}_{\text{LP}}^{[T]} / \kappa^2 - O(\text{DEV}_{\max}) && \text{(by Equation (6.11)).} \end{aligned}$$

This completes the proof of Equation (6.18).

Plugging Equation (6.18) back into Equation (6.17) and using $\kappa = 2$ we get,

$$\text{REW} \geq \frac{1}{\text{ratio}} \min \left(\frac{\text{OPT}_{\text{LP}}^{[T]}}{12d}, \frac{\text{OPT}_{\text{LP}}^{[T]}}{12} \right) - \text{reg}(T) - O(\text{DEV}_{\max}). \quad (6.19)$$

Moreover, $\text{OPT}_{\text{LP}}^{[T]} \geq \text{OPT}_{\text{FD}}$ by Equation (5.4). Plugging this into Equation (6.19) and using $\gamma \geq T^{-1/4}$, we obtain Equation (6.5), completing the proof of the theorem.

6.2 Proof of Lemma 6.9: Last Full Phase Offers Sufficient Rewards

First, we decompose the objective on a time interval as a difference between the interval's endpoints:

$$\text{OBJ}([T_1, T_2]) \geq \text{OBJ}(T_2) - \text{OBJ}(T_1 - 1). \quad (6.20)$$

This step holds for any two rounds $T_1 < T_2 \leq T$. It is proved similarly to Claim 6.8.

PROOF OF EQUATION (6.20). Let μ denote the optimal solution to $\text{LP}_{\overline{M}_{T_2, B, T_2}}$. In particular, it satisfies the consumption constraint $\mu \cdot \sum_{t \in [T_2]} c_{t,i} \leq B$. Since resource consumption is always non-negative, $\mu \cdot \sum_{t \in [T_1-1]} c_{t,i} \leq B$, which is the consumption constraint for $\text{LP}_{\overline{M}_{T_1-1, B, T_1-1}}$. So, μ is feasible for that LP as well. Consequently,

$$\sum_{t \in [T_1-1]} \mu \cdot r_t \leq \text{OBJ}(T_1 - 1).$$

Likewise μ is also feasible for $\text{LP}_{\overline{M}_{[T_1, T_2], B, [T_1, T_2]}}$, and consequently

$$\begin{aligned} \text{OBJ}([T_1, T_2]) &\geq \sum_{t \in [T_1, T_2]} \mu \cdot r_t = \sum_{t \in [T_2]} \mu \cdot r_t - \sum_{t \in [T_1-1]} \mu \cdot r_t \\ &\geq \text{OBJ}(T_2) - \text{OBJ}(T_1 - 1). \quad \square \end{aligned}$$

The rest of the proof is specific to the time interval being the last full phase.

$$\begin{aligned}
\text{OBJ}([\tau_{start}, \tau_{end}]) &\geq \text{OBJ}(\tau_{end}) - \text{OBJ}(\tau_{start} - 1) && \text{(by Equation (6.20))} \\
&\geq \text{OBJ}^{\text{ips}}(\tau_{end}) - \text{OBJ}^{\text{ips}}(\tau_{start} - 1) - 2 \cdot \text{DEV}_{\max} && \text{(by Lemma 6.5)} \\
&\geq \text{OBJ}^{\text{ips}}(\tau_{end} + 1) - \text{OBJ}^{\text{ips}}(\tau_{start}) - 2 \cdot \text{DEV}_{\max} - K/\gamma.
\end{aligned}$$

In the last inequality, we control $\text{OBJ}^{\text{ips}}(\tau_{start})$ and $\text{OBJ}^{\text{ips}}(\tau_{end})$ using, resp., Claims 6.7 and 6.8.

Let us transition from $\text{OBJ}^{\text{ips}}(\cdot)$ to guesses \widehat{g} , and use the machinery for comparing the guesses across time. For a more succinct notation, write $t = \tau_{end} + 1$. Then

$$\begin{aligned}
\text{OBJ}([\tau_{start}, \tau_{end}]) + 2 \cdot \text{DEV}_{\max} + K/\gamma &\geq \widehat{g}_t - \widehat{g}_{\tau_{start}} && \text{(by Claim 6.7)} \\
&> \widehat{g}_t \cdot (1 - 1/\kappa) && \text{(by Lemma 6.6)} \\
&\geq \widehat{g}_T \cdot (1/\kappa - 1/\kappa^2) && \text{(by Lemma 6.6)} \\
&\geq \left(\text{OPT}_{\text{LP}}^{[T]} - \text{DEV}_{\max} \right) \cdot (1/\kappa - 1/\kappa^2) && \text{(by Lemma 6.5)}.
\end{aligned}$$

Rearranging, we complete the proof of Lemma 6.9 as follows.

$$\text{OBJ}([\tau_{start}, \tau_{end}]) \geq \text{OPT}_{\text{LP}}^{[T]} \cdot (1/\kappa - 1/\kappa^2) - \text{DEV}_{\max} \cdot (1/\kappa - 1/\kappa^2 + 2) + K/\gamma.$$

6.3 Proof of Lemma 6.5 (IPS Estimators are Good)

Recall that for every $t \in [T]$ and $a \in [K]$ we have that $p_t(a)$, the probability that arm a is chosen at time t is at least $\frac{\gamma}{K}$. We now prove Lemma 6.10 which relates linear sums of rewards and consumptions computed using the unbiased estimates and the true values. Denote $R_{\gamma, \delta}(\tau) := \frac{K}{\gamma} \sqrt{2\tau \ln(T/\delta)}$.

CLAIM 6.10. *Let $\delta > 0$, $\gamma > 0$ used by the EXP3.P(γ) be given parameters. Then we have the following statements for any fixed $\mathbf{z} \in \Delta_K$.*

$$\Pr \left[\exists \tau \in [T] \left| \sum_{t \in [\tau]} \mathbf{z} \cdot [\mathbf{r}_t^{\text{ips}} - \mathbf{r}_t] \right| > R_{\gamma, \delta}(\tau) \right] \leq \delta \quad (6.21)$$

$$\forall i \in [d] \quad \Pr \left[\exists \tau \in [T] \left| \sum_{t \in [\tau]} \mathbf{z} \cdot [\mathbf{c}_{t,i}^{\text{ips}} - \mathbf{c}_{t,i}(a)] \right| > R_{\gamma, \delta}(\tau) \right] \leq \delta \quad (6.22)$$

PROOF. The proof of this follows directly from the invocation of the Azuma-Hoeffding inequality. We will show this for Equation (6.21). Define $Y_t := \mathbf{z} \cdot [\mathbf{r}_t^{\text{ips}} - \mathbf{r}_t]$ (like-wise for the lower-tail use $Y_t := \mathbf{z} \cdot (\mathbf{r}_t - \mathbf{r}_t^{\text{ips}})$). Note that this forms a martingale difference sequence since $\mathbb{E}[\mathbf{z} \cdot (\mathbf{r}_t^{\text{ips}} - \mathbf{r}_t) \mid \mathcal{H}_{t-1}] = \mathbf{z} \cdot [\mathbf{r}_t - \mathbf{r}_t] = 0$. Here we used the fact that \mathbf{z} is not random and fixed before the start of the algorithm. Also we have that $|Y_t| \leq \frac{K}{\gamma}$. Using Lemma A.1 and taking a union bound over all $\tau \in [T]$ we have the desired equation. \square

We will now prove the two inequalities in Equation (6.10). We will first prove the first inequality in Equation (6.10).

Let $\boldsymbol{\mu}^*$ denote the optimal solution to $\text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{\tau}, B(1 - \frac{R_{\gamma, \delta}(\tau)}{B}), \tau)$. Note this is valid whenever $B > \Omega(\frac{K}{\gamma} \sqrt{\tau \log \frac{T}{\delta}})$. From Equation (6.22), with probability at least $1 - \delta$ for every $i \in [d]$ we have

$$\begin{aligned}
\sum_{t \in [\tau]} \boldsymbol{\mu}^* \cdot \mathbf{c}_{t,i}^{\text{ips}} &\leq \sum_{t \in [\tau]} \boldsymbol{\mu}^* \cdot \mathbf{c}_{t,i} + R_{\gamma, \delta}(\tau). \\
&\leq B
\end{aligned} \quad (6.23)$$

Equation (6.23) used the fact that $\sum_{t \in [\tau]} \boldsymbol{\mu}^* \cdot \mathbf{c}_{t,i} \leq B(1 - R_{Y,\delta}(\tau))$. Using Equation (6.21), we have that with probability at least $1 - \delta$,

$$\sum_{t \in [\tau]} \boldsymbol{\mu}^* \cdot \mathbf{r}_t \leq \sum_{t \in [\tau]} \boldsymbol{\mu}^* \cdot \mathbf{r}_t^{\text{ips}} + R_{Y,\delta}(\tau).$$

Using the fact that,

$$\sum_{t \in [\tau]} \boldsymbol{\mu}^* \cdot \mathbf{r}_t = \text{OPT}_{\text{LP}} \left(\bar{\mathbf{M}}_\tau, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right),$$

we have the following.

$$\text{OPT}_{\text{LP}} \left(\bar{\mathbf{M}}_\tau, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right) - R_{Y,\delta}(\tau) \leq \sum_{t \in [\tau]} \boldsymbol{\mu}^* \cdot \mathbf{r}_t^{\text{ips}}. \quad (6.24)$$

From Equation (6.23) we have that $\boldsymbol{\mu}^*$ is feasible to $\text{OPT}_{\text{LP}}^{\text{ips}}(\tau)$ and from Equation (6.24) this implies that

$$\text{OBJ}^{\text{ips}}(\tau) \geq \text{OPT}_{\text{LP}} \left(\bar{\mathbf{M}}_\tau, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right) - R_{Y,\delta}(\tau). \quad (6.25)$$

Finally from Equation (6.15) we have

$$\text{OPT}_{\text{LP}} \left(\bar{\mathbf{M}}_\tau, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right) \geq \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right) \text{OBJ}(\tau). \quad (6.26)$$

From Equations (6.25) and (6.26) we have

$$\begin{aligned} \text{OBJ}^{\text{ips}}(\tau) &\geq \underbrace{\text{OBJ}(\tau) - R_{Y,\delta}(\tau) \left(1 + \frac{\text{OBJ}(\tau)}{B} \right)}_{\leq \left(1 + \frac{2\text{OBJ}(\tau)}{B} \right) \frac{K}{Y} \sqrt{2\tau \log \frac{T}{\delta}}} \\ &\leq \left(1 + \frac{2\text{OBJ}(\tau)}{B} \right) \frac{K}{Y} \sqrt{2\tau \log \frac{T}{\delta}} \end{aligned}$$

which gives the lower-tail in Equation (6.10).

We will now prove the second inequality in Equation (6.10) in a similar fashion. Let $\tilde{\boldsymbol{\mu}}^*$ denote the optimal solution to $\text{OBJ}^{\text{ips}}(\bar{\mathbf{M}}_\tau, B(1 - \frac{R_{Y,\delta}(\tau)}{B}), \tau)$.

From Equation (6.22) we have that with probability at least $1 - \delta$ for every $i \in [d]$,

$$\begin{aligned} \sum_{t \in [\tau]} \tilde{\boldsymbol{\mu}}^* \cdot \mathbf{c}_{t,i} &\leq \sum_{t \in [\tau]} \tilde{\boldsymbol{\mu}}^* \cdot \mathbf{c}_{t,i}^{\text{ips}} + R_{Y,\delta}(\tau) \\ &\leq B \end{aligned} \quad (6.27)$$

Equation (6.27) used the fact that $\sum_{t \in [\tau]} \tilde{\boldsymbol{\mu}}^* \cdot \mathbf{c}_{t,i}^{\text{ips}} \leq B(1 - R_{Y,\delta}(\tau))$. Using Equation (6.21), we have that with probability at least $1 - \delta$,

$$\sum_{t \in [\tau]} \tilde{\boldsymbol{\mu}}^* \cdot \mathbf{r}_t^{\text{ips}} \leq \sum_{t \in [\tau]} \tilde{\boldsymbol{\mu}}^* \cdot \mathbf{r}_t + R_{Y,\delta}(\tau).$$

From the fact that

$$\sum_{t \in [\tau]} \tilde{\boldsymbol{\mu}}^* \cdot \mathbf{r}_t^{\text{ips}} = \text{OPT}_{\text{LP}}^{\text{ips}} \left(\bar{\mathbf{M}}_\tau^{\text{ips}}, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right),$$

we get the following.

$$\text{OPT}_{\text{LP}}^{\text{ips}} \left(\overline{\mathbf{M}}_{\tau}^{\text{ips}}, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right) - R_{Y,\delta}(\tau) \leq \sum_{t \in [\tau]} \tilde{\boldsymbol{\mu}}^* \cdot \mathbf{r}_t. \quad (6.28)$$

From Equation (6.27) we have that $\tilde{\boldsymbol{\mu}}_j^*$ is feasible to $\text{OPT}_{\text{LP}}(\tau)$ and from Equation (6.28) this implies that

$$\text{OPT}_{\text{LP}}^{\text{ips}} \left(\overline{\mathbf{M}}_{\tau}^{\text{ips}}, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right) \leq \text{OBJ}(\tau) + R_{Y,\delta}(\tau). \quad (6.29)$$

Finally from Equation (6.15) we have

$$\text{OPT}_{\text{LP}}^{\text{ips}} \left(\overline{\mathbf{M}}_{\tau}^{\text{ips}}, B \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right), \tau \right) \geq \left(1 - \frac{R_{Y,\delta}(\tau)}{B} \right) \text{OBJ}^{\text{ips}}(\tau). \quad (6.30)$$

Combining Equations (6.30) and (6.29) we get,

$$\text{OBJ}^{\text{ips}}(\tau) \leq \text{OBJ}(\tau) + R_{Y,\delta}(\tau) + \frac{R_{Y,\delta}(\tau)}{B - R_{Y,\delta}(\tau)} (\text{OBJ}(\tau) + R_{Y,\delta}(\tau)). \quad (6.31)$$

Since $B > 2R_{Y,\delta}(\tau)$ we get,

$$\begin{aligned} \text{OBJ}^{\text{ips}}(\tau) &\leq \text{OBJ}(\tau) + \underbrace{\frac{2R_{Y,\delta}(\tau)}{B} (\text{OBJ}(\tau) + R_{Y,\delta}(\tau))}_{\leq \left(1 + \frac{2\text{OBJ}(\tau)}{B}\right)^{\frac{K}{Y}} \sqrt{2\tau \log \frac{T}{\delta}}} \end{aligned}$$

and thus we get the upper-tail in Equation (6.10).

We will now prove Equation (6.11). Recall that $\widehat{g}_{\tau} := \max_{t \in [\tau]} \text{OBJ}^{\text{ips}}(t)$. Moreover, $\text{OPT}_{\text{LP}}^{[\tau]} = \max_{t \in [\tau]} \text{OBJ}(t)$.

Consider $\widehat{g}_{\tau} - \text{OPT}_{\text{LP}}^{[\tau]}$. We have,

$$\begin{aligned} \widehat{g}_{\tau} - \text{OPT}_{\text{LP}}^{[\tau]} &= \max_{t \in [\tau]} \text{OBJ}^{\text{ips}}(t) - \text{OPT}_{\text{LP}}^{[\tau]} \\ &\leq \max_{t \in [\tau]} (\text{OBJ}(t) + \text{DEV}(t)) - \text{OPT}_{\text{LP}}^{[\tau]} \\ &\leq \max_{t \in [\tau]} \text{OBJ}(t) + \max_{t \in [\tau]} \text{DEV}(t) - \text{OPT}_{\text{LP}}^{[\tau]} \\ &= \max_{t \in [\tau]} \text{DEV}(t). \end{aligned}$$

Now consider $\text{OPT}_{\text{LP}}^{[\tau]} - \widehat{g}_{\tau}$. We have,

$$\begin{aligned} \text{OPT}_{\text{LP}}^{[\tau]} - \widehat{g}_{\tau} &\leq \text{OPT}_{\text{LP}}^{[\tau]} - \max_{t \in [\tau]} (\text{OBJ}(t) - \text{DEV}(t)) \\ &\leq \text{OPT}_{\text{LP}}^{[\tau]} - \max_{t \in [\tau]} \text{OBJ}(t) + \max_{t \in [\tau]} \text{DEV}(t) \\ &= \max_{t \in [\tau]} \text{DEV}(t). \end{aligned}$$

This completes the proof of Lemma 6.5.

7 EXTENSIONS

We obtain several extensions which highlight the modularity of LagrangeBwK: we apply Theorems 4.4 and 5.1 with appropriately chosen primal algorithm ALG_1 , and immediately obtain strong performance guarantees.¹³ We tackle four well-known scenarios:

- *full feedback* [e.g., 10, 52, 70]: in each round, the algorithm chooses an action and observes the outcomes of all possible actions; this is a classic scenario in online machine learning.
- *combinatorial semi-bandits* [e.g., 12, 59, 64]: actions are feasible subsets of “atoms”. The atoms in the chosen action have individual outcomes that are observed and add up to the action’s total outcome. Typical motivating example are subsets/lists of news articles, ads, or web search results.
- *contextual bandits with policy sets* [e.g., 4, 47, 68]: before each round, a *context* is observed, and the algorithm competes against the best policy (mapping from context to actions) in a given policy class. In a typical application scenario, the context includes known features of the current user.
- *bandit convex optimization* (starting from Flaxman et al. [50], Kleinberg [67], with recent advances Bubeck et al. [29, 30], Hazan and Levy [61]). Here the set of actions is a convex set $\mathcal{X} \subset \mathbb{R}^K$. For each round t , the adversary chooses a concave function $f_t : \mathcal{X} \rightarrow [0, 1]$ such that the reward for chosen action $\mathbf{x} \in \mathcal{X}$ is $f_t(\mathbf{x})$.

Formalities. To simplify the statements, we make the following assumptions without further mention:

- The dual algorithm, ALG_2 , is always Hedge, with the associated regret bound from Equation (3.4). For high-probability regret bounds, $\delta = \frac{1}{T}$ is a fixed and known failure probability parameter.
- For Stochastic BwK, one resource is the dummy resource (with consumption $\frac{B}{T}$ for each arm). Algorithm LagrangeBwK is run with parameters $B_0 = B$ and $T_0 = T$.
- For Adversarial BwK, one of the arms is a *null arm* that has zero reward and zero resource consumption. Algorithm 2 is run with any $\kappa > 1$ and range $[g_{\min}, g_{\max}] = [\sqrt{T}, T]$, as in Theorem 5.1(b).

A typical corollary. All our corollaries have the following shape, for some regret term reg :

(C1) In the stochastic version, algorithm LagrangeBwK achieves, with probability at least $1 - \frac{1}{T}$,

$$\text{OPT}_{\text{DP}} - \text{REW} \leq O\left(\frac{T}{B} \cdot \text{reg}\right).$$

(C2) In the adversarial version, Algorithm 2 achieves

$$\frac{\mathbb{E}[\text{REW}]}{\text{OPT}_{\text{FD}}} \geq \frac{2}{(d+1) \ln(T)} - O(\text{reg}) \left(\frac{1}{\text{OPT}_{\text{FD}}} + \frac{1}{dB} \right).$$

Corollaries similar to (C2) can be achieved for Algorithm 3, too; we omit them for ease of exposition.

¹³For these theorems to hold, ALG_1 needs to satisfy regret bound (3.2) only against adaptive adversaries that arise in the repeated Lagrange game in the corresponding extension, not against *arbitrary* adaptive adversaries.

7.1 BwK with Full Feedback

In the full-feedback version of BwK, the entire outcome matrix \mathbf{M}_t is revealed to the algorithm after each round t . Accordingly, we can use Hedge as the primal algorithm ALG_1 . The effect is, essentially, that the dependence on K , the number of arms, in the regret term becomes logarithmic rather than \sqrt{K} .

COROLLARY 7.1. *Consider BwK with full feedback. Using Hedge as the primal algorithm, we obtain corollaries (C1) and (C2) with regret term $\text{reg} = \sqrt{T \ln(dKT)}$.*

Adversarial BwK with full feedback have not been studied before. For the stochastic version, the regret bound is unsurprising: one expects to obtain a similar improvement with each of the three other algorithms in the prior work on Stochastic BwK by tracing the “confidence terms” through the analysis. The significance here is that we obtain this result as an immediate corollary.

7.2 Combinatorial Semi-Bandits with Knapsacks

Following [86], we consider *Combinatorial Semi-BwK*, a common generalization of BwK and *combinatorial semi-bandits* [e.g., 12, 59, 64]. In this problem, actions correspond to subsets of some finite ground set Ω of size n , whose elements are called *atoms*. There is a fixed family $\mathcal{F} \subset 2^\Omega$ of feasible actions. For each round t , there is an outcome vector $\mathbf{o}_{t,e} \in [0, \frac{1}{n}]^{d+1}$ for each round atom $e \in \Omega$, with the same semantics as the actions’ outcome vectors. If an action $S \subset \Omega$ is chosen, the outcome vectors $\mathbf{o}_{t,e}$ are observed for all atoms $e \in S$, and the action’s outcome is the sum $\mathbf{M}_t(S) = \sum_{e \in S} \mathbf{o}_{t,e} \in [0, 1]^d$. In the adversarial case, all outcome vectors $\mathbf{o}_{t,e}$, $t \in [T]$, $e \in \Omega$ are chosen by an adversary arbitrarily before round 1. In the stochastic case, the atomic outcome matrix $(\mathbf{o}_{t,e} : e \in \Omega)$ is drawn independently in each round t from some fixed distribution. Combinatorial semi-bandits, as studied previously, is a special case with no resource constraints ($d = 0$).

Typical motivating examples involve ad/content personalization scenarios. Atoms can correspond to items news articles, ads, or web search results, and actions are subsets that satisfy some constraints on quantity, relevance, or diversity of items. One can also model ranked lists of atoms: then atoms are rank-item pairs, and each feasible action $S \subset \Omega$ satisfies a constraint that each rank between 1 and $|S|$ is present in exactly one chosen rank-item pair.

A naive application of our main results suffers from regret terms that are proportional to $\sqrt{|\mathcal{F}|}$, which may be exponential in the number of atoms n . Instead, the work on combinatorial semi-bandits features regret bounds that scale polynomially in n . This is what we achieve, too. We use an algorithm from [79] which solves combinatorial semi-bandits in the absence of resource constraints. This algorithm satisfies a high-probability regret bound (3.2) against an adaptive adversary, with $R_\delta(T) = O(\sqrt{nT \log(1/\delta)})$.¹⁴

COROLLARY 7.2. *Consider Combinatorial Semi-BwK with n atoms. Using the algorithm from [79] as the primal algorithm, we obtain corollaries (C1) and (C2) with regret term $\text{reg} = \sqrt{nT \log T}$.*

The adversarial version of Combinatorial Semi-BwK has not been studied before.

The stochastic version has been studied in [86] when the action set is a matroid, achieving regret

$$\tilde{O}\left(\text{OPT}_{\text{DP}} \sqrt{n/B} + \sqrt{T/n} + \sqrt{\text{OPT}_{\text{DP}}}\right).$$

This regret bound becomes $\tilde{O}(\sqrt{nT})$ in the regime when B and OPT_{DP} are $\Omega(T)$ (see Footnote 14). We achieve the same regret bound for this regime, without the matroid assumption and without

¹⁴Prior work [79, 86] posits that atoms’ per-round rewards/consumptions lie in the range $[0, 1]$, rather than $[0, \frac{1}{n}]$, so their stated regret bounds should be recaled accordingly.

any extra work. However, the regret bound in [86] can be substantially better than ours when $\text{OPT}_{\text{DP}} \ll T$.

7.3 Contextual Bandits with Knapsacks

Following [7, 20], we consider **Contextual Bandits with Knapsacks (cBwK)**, a common generalization of BwK and *contextual bandits with policy sets* [e.g., 4, 47, 68]. The only change in the protocol, compared to BwK, is that in the beginning of each round t a context $x_t \in \mathcal{X}$ arrives and is observed by the algorithm before it chooses an action. The context set \mathcal{X} is arbitrary and known. In the **adversarial version (Adversarial cBwK)** both x_t and the outcome matrix \mathbf{M}_t is chosen by an adversary. In the **stochastic version (Stochastic cBwK)** the pair (x_t, \mathbf{M}_t) is chosen independently from some fixed but unknown distribution over such pairs.

In cBwK one is also given a finite set Π of *policies*: deterministic mappings from contexts to actions.¹⁵ Essentially, the algorithm competes with the best course of action restricted to these policies. For a formal definition, let us interpret cBwK as a BwK problem with action set Π , denote this problem as $\text{BwK}(\Pi)$. In other words, actions in $\text{BwK}(\Pi)$ are policies in cBwK. An algorithm for $\text{BwK}(\Pi)$ is oblivious to context arrivals. It chooses a policy $\pi_t \in \Pi$ in each round t , and receives an outcome for this policy: namely, the outcome for action $\pi(x_t)$. We are interested in the usual benchmarks for this problem, the best algorithm OPT_{DP} and the best fixed distribution OPT_{FD} (where both benchmarks are constrained to use policies in Π); denote them $\text{OPT}_{\text{DP}}(\Pi)$ and $\text{OPT}_{\text{FD}}(\Pi)$, respectively.

Without budget constraints (i.e., with $B = T$), this is exactly contextual bandits with policy set Π . Both benchmarks then reduce to the standard benchmark of the best fixed policy.

Background: algorithm EXP4.P. We use EXP4.P [26], an algorithm for the contextual version of adversarial online learning with bandit feedback. The algorithm operates according to the protocol in Figure 2.

We are interested in regret bounds for EXP4.P relative to the best fixed policy:

$$\text{OPT}_{\Pi} = \max_{\pi \in \Pi} \sum_{t \in [T]} f_t(\pi(x_t)).$$

For each round t , the pair (x_t, \mathbf{g}_t) induces a payoff vector $\mathbf{f}_t \in [b_{\min}, b_{\max}]^{\Pi}$ on policies:

$$\mathbf{f}_t(\pi) = g_t(\pi(x_t)) \quad \forall \pi \in \Pi.$$

THEOREM 7.3 ([26]). *Fix failure probability $\delta > 0$, policy set Π , and payoff range $[b_{\min}, b_{\max}]$. Then algorithm EXP4.P (appropriately tuned) satisfies the following regret bound:*

$$\Pr \left[\text{OPT}_{\Pi} - \sum_{t \in [T]} f_t(\pi_t) \leq (b_{\max} - b_{\min}) R_{\delta}(T) \right] \geq 1 - \delta, \quad (7.1)$$

with regret term $R_{\delta}(T) = O\left(\sqrt{\tau K \log(KT |\Pi|/\delta)}\right)$.

Our solution for cBwK. We solve cBwK by reducing it to $\text{BwK}(\Pi)$, and treating it as a BwK problem. A naive solution simply posits $|\Pi|$ arms and directly applies the machinery developed earlier in this paper. This results in $\sqrt{|\Pi|}$ dependence in regret bounds, which is unsatisfactory, as the policy set may be very large. Instead, we use EXP4.P as the primal algorithm (ALG_1). We interpret EXP4.P as an algorithm for (non-contextual) adversarial online learning, as defined in Section 3,

¹⁵W.l.o.g. assume that Π contains all *constant policies*, i.e., all policies that always evaluate to the same action.

- Given:** action set $[K]$, context set \mathcal{X} , policy set Π , payoff range $[b_{\min}, b_{\max}]$.
 In each round $t \in [T]$,
1. the adversary chooses a context $x_t \in \mathcal{X}$ and a payoff vector $\mathbf{g}_t \in [b_{\min}, b_{\max}]^K$;
 2. the algorithm chooses a distribution \mathbf{p}_t over Π (without seeing x_t);
 3. a policy $\pi_t \in \Pi$ is drawn independently from \mathbf{p}_t ;
 4. algorithm's chosen action is defined as $a_t = \pi_t(x_t) \in [K]$;
 5. payoff $g_t(a_t)$ is received by the algorithm.

Fig. 2. Adversarial contextual bandits.

with action set Π . It is easy to see that Theorem 7.3 provides regret bound (3.2) under this interpretation. Therefore, we obtain the following:

COROLLARY 7.4. *Consider contextual bandits with knapsacks, with policy set Π . Using EXP4.P as the primal algorithm, we obtain corollaries (C1) and (C2) with regret term $\text{reg} = \sqrt{TK \ln(dKT |\Pi|)}$. The benchmarks are $\text{OPT}_{\text{DP}} = \text{OPT}_{\text{DP}}(\Pi)$ and $\text{OPT}_{\text{FD}} = \text{OPT}_{\text{FD}}(\Pi)$.*

Adversarial cBwK has not been studied before. The regret bound for the adversarial case is meaningful only if $B > \sqrt{T}$. This is essentially inevitable in light of the lower bound in Theorem 8.3.

Stochastic cBwK has been studied in [7, 20], achieving regret

$$O(\text{reg})(1 + \text{OPT}_{\text{DP}}(\Pi)/B), \quad (7.2)$$

where the reg term is the same as in Corollary 7.4. Whereas the regret bound from Corollary 7.4 is $O(\text{reg} \cdot T/B)$. Note that we match (7.2) in the regime $\text{OPT}_{\text{DP}}(\Pi) > \Omega(T)$. Our regret bound is optimal, up to logarithmic factors, in the regime $B > \Omega(T)$. This is due to the $\min\left(T, \Omega(\sqrt{KT \log(|\Pi|)/\log(K)})\right)$ lower bound on regret, which holds for contextual bandits [3].

Discussion. Our algorithms are slow, as the per-round running time of EXP4.P is proportional to $|\Pi|$. The state-of-art approach to computational efficiency in contextual bandits is *oracle-efficient algorithms*, which make only a small number of calls to an oracle that finds the best policy in Π for a given data set. In particular, prior work for Stochastic cBwK [7] obtains an oracle-efficient algorithm with regret bound as in (7.2). To obtain oracle-efficient algorithms for cBwK in our framework, both for the stochastic and adversarial versions, it suffices to replace EXP4.P with an oracle-efficient algorithm for adversarial contextual bandits that obtains regret bound (7.1), possibly with a larger regret term R_δ . Such algorithms *almost* exist: a recent breakthrough [81, 93, 94] obtains algorithms with similar regret bounds, but only for expected regret.

7.4 Bandit Convex Optimization with Knapsacks

We consider **Bandit Convex Optimization with Knapsacks (BCOwK)**, a common generalization of BwK and *bandit convex optimization*. We define BCOwK as a version of BwK, where the action set \mathcal{X} is a convex subset of \mathbb{R}^K . For each round t , there is a concave function $f_t : \mathcal{X} \rightarrow [0, 1]$ and convex functions $g_{t,i} : \mathcal{X} \rightarrow [0, 1]$, for each resource i , so that the reward for choosing action $\mathbf{x} \in \mathcal{X}$ in this round is $f_t(\mathbf{x})$ and consumption of each resource i is $g_{t,i}(\mathbf{x})$. In the stochastic version, the tuple of functions $(f_t; g_{t,1}, \dots, g_{t,d})$ is sampled independently in each round t from some fixed distribution (which is not known to the algorithm). In the adversarial version, all these tuples are chosen by an adversary before the first round.

Neither the stochastic nor adversarial version of BCOwK have been studied in prior work (but see the discussion of constrained online convex optimization in Section 2). Bandit convex optimization, as studied previously, is a special case with no resource constraints ($d = 0$).

The primal algorithm ALG_1 in LagrangeBwK faces an instance of BCO (with an adaptive adversary). This is because the Lagrange function (4.3) is a concave function of the action, as sum of concave functions. For our primal algorithm, we use a recent breakthrough on BCO due to [30]. This algorithm satisfies the high-probability regret bound (3.2) against an adaptive adversary, with regret term

$$R_\delta(T) = O(K^{9.5} \log^7(T) \sqrt{T \log(1/\delta)}).$$

We assume the existence of a *null arm*: a point $\mathbf{x} \in \mathcal{X}$ such that $f_t(\mathbf{x}) = g_{t,i}(\mathbf{x}) = 0$ for each resource i except the “dummy resource”. (Recall that we posit the “dummy resource” – a resource whose consumption is B/T for each arm – for the stochastic version.) Unlike elsewhere in this paper, this assumption is not without loss of generality: indeed, the null arm should be “embedded” into \mathcal{X} without breaking the convexity/concavity properties. Moreover, we assume that the null arm lies in the interior of \mathcal{X} .

COROLLARY 7.5. *Consider BCOwK for a given convex set $\mathcal{X} \subset \mathbb{R}^K$. Using the algorithm from [30] as the primal algorithm, we obtain corollaries (C1) and (C2) with regret term $\text{reg} = K^{9.5} \log^{7.5}(T) \sqrt{T}$.*

Remark 7.6. LagrangeBwK framework extends to infinite action sets: everything carries over, as long as Equation (4.4) holds. (Essentially, we never take union bounds over actions, and we can replace max and sums over actions with sup and integrals.) For BCOwK , Equation (4.4) is a statement about constrained convex optimization programs. According to *Slater’s condition* [see Equation (5.27) in 27], it suffices to have a point \mathbf{x} in the interior of \mathcal{X} such that $g_{t,i}(\mathbf{x}) < B/T$ for each resource $i \in [d]$ other than the dummy resource (or any other resource whose consumption is the same in all rounds). One such point is the null arm.

8 LOWER BOUNDS

We prove the lower bounds on the competitive ratio that we have claimed in Section 1: the $\Omega(\log T)$ lower bound w.r.t. the best fixed distribution benchmark (OPT_{FD}), the $\Omega(T)$ lower bound w.r.t. the best dynamic policy benchmark (OPT_{DP}), and the $\Omega(K)$ lower bound w.r.t. the best fixed arm benchmark (OPT_{FA}). As a warm-up, we analyze the simple example from Section 1 that leads to the $\frac{5}{4}$ lower bound w.r.t. OPT_{FD} . All lower-bounds are for a randomized algorithm against an oblivious adversary. We summarize all these lower bounds in the following theorem:

THEOREM 8.1. *Consider Adversarial BwK with a single resource ($d = 1$), K arms, budget B , and time horizon T . Consider any randomized algorithm for this problem, and let REW denote its reward. Then:*

- (a) $\text{OPT}_{\text{FD}}/\mathbb{E}[\text{REW}] \geq \frac{5}{4} - o(1)$ for some problem instance (from the example in the Introduction).
This holds even if $\text{OPT}_{\text{FD}} \geq T/4$ and $B = T/2$.
- (b) $\text{OPT}_{\text{FD}}/\mathbb{E}[\text{REW}] \geq \Omega(\log T)$ for some problem instance with $K = 2$ arms.
This holds for any given budget $B \in [c_0 \log^3(T), O(T^{1-\alpha})]$, even if $\text{OPT}_{\text{FD}} \geq B^2/T$.
Here $\alpha \in (0, 1)$ is an arbitrary absolute constant, and c_0 is any large enough absolute constant.
- (c) $\text{OPT}_{\text{DP}}/\mathbb{E}[\text{REW}] \geq T/B^2$ for some problem instance with $K = 2$ arms.
This holds for any given budget $B < \sqrt{T}$, with $\text{OPT}_{\text{FD}} = B$.
- (d) $\text{OPT}_{\text{FA}}/\mathbb{E}[\text{REW}] \geq \Omega(K)$ for some problem instance.
This holds for any given budget B , with $\text{OPT}_{\text{FA}} > B/K^K$.

Remark 8.2. The lower bounds for parts (a,b,c) hold (even) for problem instances with $K = 2$ arms, one of which is the “null arm” with no rewards and no resource consumption. The lower bounds in parts (a,b) hold even for a much more permissive feedback model from the online packing literature, namely, when the algorithm observes the outcome vector for all actions in a given round, and moreover does it *before* it chooses an arm in this round.

We tweak our construction from Theorem 8.1(c) to obtain a strong lower bound for the contextual version of Adversarial BwK (a.k.a. *Adversarial cBwK*), as studied in Section 7.3. This lower bound implies that Adversarial cBwK is essentially hopeless in the regime $B < \sqrt{T}$, complementing a strong positive result (Corollary 7.4) for the regime $B > \tilde{\Omega}(\sqrt{T})$. It is proved in Section 8.3, along with Theorem 8.1(c).

THEOREM 8.3. *Consider adversarial contextual bandits with knapsacks (Adversarial cBwK), with policy class Π , a single resource ($d = 1$), $K = 2$ arms, and any given budget $B < \sqrt{T}$. Consider any randomized algorithm for this problem, and let REW denote its reward. Then*

$$\text{OPT}_{\text{FD}}(\Pi)/\mathbb{E}[\text{REW}] \geq T/B^2 \quad \text{for some problem instance.}$$

Notation. In the proof of lower-bounds below, we use the following notation. Given an instance \mathcal{I} , we denote $\text{OPT}_{\text{FD}}(\mathcal{I})$, $\text{OPT}_{\text{FA}}(\mathcal{I})$ and $\text{OPT}_{\text{DP}}(\mathcal{I})$ to denote the optimal value of the best fixed distribution, best fixed arm and best dynamic policy respectively, for instance \mathcal{I} . Likewise let $\text{OPT}_{\text{LP}}^{[T]}(\mathcal{I})$ denote the optimal LP value for instance \mathcal{I} and given an algorithm \mathcal{A} and an instance \mathcal{I} , let $\mathbb{E}[\text{REW}(\mathcal{A}, \mathcal{I})]$ denote the expected reward obtained by \mathcal{A} on instance \mathcal{I} , where the expectation is over the internal randomness of the algorithm.

8.1 Warm-up: Example from the Introduction

As a warm-up, let us recap and analyze the example from the Introduction.

CONSTRUCTION 8.4. *There are two arms and one resource with budget $B = \frac{T}{2}$. Arm 1 has zero rewards and zero consumption. Arm 2 has consumption 1 in each round, and offers reward $\frac{1}{2}$ in each round of the first half-time ($\frac{T}{2}$ rounds). In the second half-time, arm 2 offers either reward 1 in all rounds, or reward 0 in all rounds. More formally, there are two problem instances, call them \mathcal{I}_1 and \mathcal{I}_2 , that coincide for the first half-time and differ in the second half-time.*

LEMMA 8.5. *Any algorithm suffers $\text{OPT}_{\text{FD}}/\mathbb{E}[\text{REW}] \geq \frac{5}{4} - o(1)$ on some instance in Construction 8.4.*

The intuition is that given a random instance as input the algorithm needs to choose how much budget to invest in the first half-time, without knowing what comes in the second, and any choice (in expectation) leads to the claimed competitive ratio.

To prove Lemma 8.5 (as well as the main lower bound in Theorem 8.1(b)), we compare the algorithm's performance to $\text{OPT}_{\text{LP}}^{[T]}$, and invoke the following lemma:

LEMMA 8.6. $\text{OPT}_{\text{FD}} \geq \text{OPT}_{\text{LP}}^{[T]} - O(\text{OPT}_{\text{LP}}^{[T]} \cdot \sqrt{\frac{\log dT}{B}})$.

PROOF. Let τ^* denote the time-step at which $\text{OPT}_{\text{LP}}^{[T]}$ is maximized. Let \mathbf{p} denote the optimal solution to $\tau^* \cdot \text{OPT}_{\text{LP}}(\bar{\mathbf{M}}_{\tau^*}, B(1 - \epsilon), \tau^*)$ where $\epsilon = \sqrt{\frac{\log dT}{B}}$. Note that OPT_{FD} is at least as large as the expected total reward obtained by the distribution \mathbf{p} . From the Chernoff-Hoeffding bounds (Lemma A.2), with probability at least $1 - dT^{-2}$ we have

$$\forall i \in [d] \quad \sum_{t \in [\tau^*]} \mathbf{p} \cdot \mathbf{c}_{t,i} \leq B.$$

Conditioning on this event the expected total reward obtained by \mathbf{p} is

$$\sum_{t \in [\tau^*]} \mathbf{p} \cdot \mathbf{r}_t = \tau^* \cdot \text{OPT}_{\text{LP}}(\bar{\mathbf{M}}_{\tau^*}, B(1 - \epsilon), \tau^*).$$

Thus, the expected total reward obtained by \mathbf{p} is at least $\tau^* \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{\tau^*}, B(1 - \epsilon), \tau^*)$.¹⁶ Moreover, from Equation (6.15) we have that

$$\begin{aligned} \text{OPT}_{\text{FD}} &\geq \tau^* \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{\tau^*}, B(1 - \epsilon), \tau^*) \\ &\geq (1 - \epsilon)\tau^* \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{\tau^*}, B, \tau^*) \\ &\geq \text{OPT}_{\text{LP}}^{[T]} - O\left(\text{OPT}_{\text{LP}}^{[T]} \sqrt{\frac{\log dT}{B}}\right). \end{aligned} \quad \square$$

PROOF OF LEMMA 8.5. Denote the two arms by A_1 and A_0 where A_0 denotes the null arm. The consumption for arm A_1 is 1 for all rounds in both \mathcal{I}_1 and \mathcal{I}_2 . Thus, the only difference between the two instances is the rewards obtained for playing arm A_1 in each round. The instances have two phases where each phase lasts for $\frac{T}{2}$ rounds. In phase 1, in both \mathcal{I}_1 and \mathcal{I}_2 playing arm A_1 fetches a reward $\frac{1}{2}$. In the second phase, in \mathcal{I}_1 , the reward for playing arm A_1 is 0 while in \mathcal{I}_2 the reward for playing arm A_1 is 1. Thus, the *outcome* matrix for the first $\frac{T}{2}$ time-steps is the same in instances \mathcal{I}_1 and \mathcal{I}_2 .

Consider a randomized algorithm \mathcal{A} . Let α_1 be the expected number of times arm A_1 is played by \mathcal{A} in the first $\frac{T}{2}$ rounds on instances \mathcal{I}_1 and \mathcal{I}_2 . Note since the outcome matrix is the same, the expected number of times the arm is played should be same in both the instances. Let $\alpha_{2,1}$, $\alpha_{2,2}$ denote the expected number of times arm A_1 is played in the second phase in instances \mathcal{I}_1 and \mathcal{I}_2 , respectively.

Recall that in this section we are interested in a lower-bound on the competitive ratio $\text{OPT}_{\text{FD}}/\mathbb{E}[\text{REW}]$ for every instance. Consider $\text{OPT}_{\text{LP}}^{[T]}(\mathcal{I}_1)$, the optimal value of the best fixed distribution on \mathcal{I}_1 . Using Equation (5.4) with $\tau = \frac{T}{2}$ this equals $\frac{T}{2} \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_{\frac{T}{2}, B, \frac{T}{2}})$ which evaluates to $\frac{T}{4}$. Likewise $\text{OPT}_{\text{LP}}^{[T]}(\mathcal{I}_2)$ equals $T \cdot \text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_T, B, T)$, which evaluates to $\frac{3T}{8}$. Consider the performance of \mathcal{A} on \mathcal{I}_1 . We have,

$$\frac{\text{OPT}_{\text{LP}}^{[T]}(\mathcal{I}_1)}{\mathbb{E}[\text{REW}(\mathcal{A}, \mathcal{I}_1)]} \geq \left(\frac{T}{4}\right) \bigg/ \left(\frac{\alpha_1}{2}\right). \quad (8.1)$$

Likewise on \mathcal{I}_2 we have,

$$\frac{\text{OPT}_{\text{LP}}^{[T]}(\mathcal{I}_2)}{\mathbb{E}[\text{REW}(\mathcal{A}, \mathcal{I}_2)]} \geq \left(\frac{3T}{8}\right) \bigg/ \left(\frac{\alpha_1}{2} + \alpha_{2,2}\right). \quad (8.2)$$

Thus, the competitive ratio of \mathcal{A} is at least the maximum of the ratios in Equations (8.1) and (8.2). Thus, we want to minimize this maximum and is achieved when the two ratios are equal to each other.

Notice that the term $\alpha_{2,1}$ does not appear in Equations (8.1) and (8.2). By setting the term in Equation (8.1) equal to the term in Equation (8.2) and re-arranging,

$$\alpha_1 = 4\alpha_{2,2}. \quad (8.3)$$

Moreover, we have $\alpha_1 + \alpha_{2,2} \leq B$. Combining this with Equation (8.3) we get $\alpha_1 \leq \frac{4B}{5} = \frac{2T}{5}$ and the corresponding competitive ratio is at least $(\frac{T}{4})/(\frac{\alpha_1}{2}) \geq \frac{5}{4}$. By Lemma 8.6 with $d = 1$, for every $j \in [2]$,

$$\text{OPT}_{\text{FD}}(\mathcal{I}_j) / \mathbb{E}[\text{REW}(\mathcal{A}, \mathcal{I}_j)] \geq \frac{5}{4} - O\left(\frac{\text{OPT}_{\text{LP}}^{[T]}}{T} \sqrt{\frac{\log T}{B}}\right). \quad \square$$

¹⁶With probability T^{-2} we assume that \mathbf{p} has an expected reward of 0.




Fig. 3. The lower-bounding construction for the $\Omega(\log T)$ lower bound. Here $\epsilon = \frac{B}{T}$.

8.2 The Main Lower Bound: Proof of Theorem 8.1(b)

To obtain the $\Omega(\log T)$ lower bound in Theorem 8.1(b), we extend Construction 8.4 to one with $\Omega(\log T)$ phases rather than just two. As before, the algorithm needs to decide how much budget to save for the subsequent phases; without knowing whether they would bring higher rewards or nothing. The construction is as follows, see Figure 3 for a pictorial representation:

CONSTRUCTION 8.7. *There is one resource with budget B , and two arms, denoted A_0, A_1 . Arm A_0 is the “null arm” that has zero reward and zero consumption. The consumption of arm A_1 is 1 in all rounds. The rewards of A_1 are defined as follows. We partition the time into $\frac{T}{B}$ phases of duration B each (for simplicity, assume that B divides T). We consider $\frac{T}{B}$ problem instances; for each instance \mathcal{I}_τ , $\tau \in [\frac{T}{B}]$ arm A_1 has positive rewards up to and including phase τ ; after that all rewards are 0. In each phase $\sigma \in [\tau]$, arm A_1 has reward $\sigma B/T$ in each round.*

The lower bound holds for a wide range of budgets B , as expressed by the following lemma:

LEMMA 8.8. *For any budget B and any algorithm there is a problem instance in Construction 8.7 such that*

$$\frac{\text{OPT}_{\text{FD}}}{\mathbb{E}[\text{REW}]} \geq \frac{1}{2} \cdot \ln(\lfloor T/B \rfloor) + \zeta - O\left(\frac{\log^{1.5} T}{\sqrt{B}}\right), \quad (8.4)$$

where $\zeta = 0.577\dots$ is the Euler-Masceroni constant, and $\text{OPT}_{\text{FD}} \geq B^2/T$.

In the rest of this subsection we prove Lemma 8.8. Fix any randomized algorithm \mathcal{A} . As before in this sub-section we are interested in the ratio $\text{OPT}_{\text{FD}}/\mathbb{E}[\text{REW}(\mathcal{A})]$. We argue that it has the claimed competitive ratio on at least one instance \mathcal{I}_τ in the construction 8.7. The proof proceeds in two parts. We first argue about the solution structure of the optimal distribution for the construction 8.7 (we prove this formally in Lemma 8.9). Next we characterize the expected number of times arm A_1 is played if \mathcal{A} optimal algorithm in each of the phases. Combining the two we get Lemma 8.8.

Structure of the optimal solution. Define $\text{OPT}_{\text{LP}}(\overline{M}_{\tau^*}, B, \tau^*)$ to be the optimal value of LP 4.1 on the instance \mathcal{I}_τ . Then we have the following Lemma.

LEMMA 8.9. *For a given instance \mathcal{I}_τ we have $\text{OPT}_{\text{LP}}(\overline{M}_{\tau^*}, B, \tau^*) = \frac{\epsilon B(\tau+1)}{2}$.*

PROOF. Let $\mathcal{P}(t)$ denote the non-zero reward on arm A_1 at time-step t (i.e., $\mathcal{P}(t) = \lceil \frac{t}{B} \rceil \epsilon$). It suffices to prove that the optimal stopping time $\tau^* = B\tau$. Indeed, given that the stopping time is $B\tau$, the optimal solution is to set $X(1) = \frac{1}{\tau}$ and $X(0) = 1 - \frac{1}{\tau}$ thus obtaining a total reward of $\frac{1}{\tau} \sum_{t \in [B\tau]} \mathcal{P}(t)\epsilon$. From the definition of $\mathcal{P}(t)$ we have that $\frac{1}{\tau} \sum_{t \in [B\tau]} \mathcal{P}(t)\epsilon = \frac{1}{\tau} \sum_{j \in [\tau]} \epsilon B j$. Using the fact that $\sum_{j \in [\tau]} j = \frac{\tau(\tau+1)}{2}$ we get the statement of the Lemma. Thus, it remains to prove that the optimal stopping time $\tau^* = B\tau$.

First it is easy to prove that $\tau^* \leq B\tau$. Since there are no rewards after time-step τ^* , we have

$$\forall t' > 0 \quad \text{OPT}_{\text{LP}}(\overline{\mathcal{M}}_{\tau^*+t'}, B, \tau^* + t') = \frac{1}{\tau + t'} \sum_{t \in [\tau^*]} \mathcal{P}(t)\epsilon < \frac{1}{\tau} \sum_{t \in [B\tau]} \mathcal{P}(t)\epsilon.$$

Now we will argue that the optimal stopping time cannot be strictly lesser than τ^* . To do so, first we argue that for two stopping times $t_1 < t_2$ within the same phase, the maximum objective is achieved for the stopping time t_2 . This implies that the optimal stopping time has to be the last time step of some phase.

Consider times $t_1 < t_2$ such that $\mathcal{P}(t_1) = \mathcal{P}(t_2) = \tau$. Then we want to claim that

$$\frac{B}{t_1} \left(\sum_{t \in [t_1]} \mathcal{P}(t)\epsilon \right) \leq \frac{B}{t_2} \left(\sum_{t \in [t_2]} \mathcal{P}(t)\epsilon \right).$$

For contradiction assume the inequality does not hold. Then we have the following.

$$\sum_{t \in [t_1]} \mathcal{P}(t) > \frac{t_1}{t_2} \left(\sum_{t \in [t_2]} \mathcal{P}(t) \right).$$

Note that $\sum_{t \in [B(\tau-1)]} \mathcal{P}(t) = \sum_{t' \in [\tau]} Bt' = \frac{B(\tau-1)\tau}{2}$. Thus we have

$$\begin{aligned} \sum_{t \in [t_1]} \mathcal{P}(t) &= \frac{B(\tau-1)\tau}{2} + (t_1 - B(\tau-1))\tau, \\ \sum_{t \in [t_2]} \mathcal{P}(t) &= \frac{B(\tau-1)\tau}{2} + (t_2 - B(\tau-1))\tau. \end{aligned}$$

Therefore we have,

$$\frac{B(\tau-1)\tau}{2} + (t_1 - B(\tau-1))\tau > \frac{t_1 B(\tau-1)\tau}{2t_2} + \frac{t_1}{t_2} \cdot (t_2 - B(\tau-1))\tau.$$

Further re-arranging, we get $B(\tau-1) > t_1$. This is a contradiction since t_1 is in phase τ , so $t_1 \geq B(\tau-1)$.

Next we argue that the optimal value is achieved when the stopping time is in the last *non-zero rewards phase*. Consider two phases $\tau_1 < \tau_2$. Thus, the ending times are $B\tau_1$ and $B\tau_2$. To prove that the optimal value increases by stopping at $B\tau_2$, as opposed to $B\tau_1$, we want to show that

$$\frac{1}{\tau_1} \sum_{t \in [\tau_1]} Bt\epsilon \leq \frac{1}{\tau_2} \sum_{t \in [\tau_2]} Bt\epsilon.$$

As before assume for a contradiction that this doesn't hold. Then re-arranging we get, $\frac{\tau_1(\tau_1+1)}{2} > \frac{\tau_1(\tau_2+1)}{2}$, which implies $\tau_1 > \tau_2$, contradiction. We conclude that the stopping time is $\tau^* = B\tau$. \square

Expected behavior of the optimal algorithm. Consider any randomized algorithm \mathcal{A} . The performance of \mathcal{A} is then as follows. From the definition of $\text{OPT}_{\text{LP}}^{[T]}$ we have,

$$\frac{\text{OPT}_{\text{LP}}^{[T]}}{\mathbb{E}[\text{REW}(\mathcal{A})]} = \max_{1 \leq \tau \leq T/B} \frac{B\tau \cdot \text{OPT}_{\text{LP}}(\overline{\mathcal{M}}_{B\tau}, B, B\tau)}{\mathbb{E}[\text{REW}(\mathcal{A})]}. \quad (8.5)$$

We will now show that for any algorithm \mathcal{A} , there exists an instance $j \in \left[\frac{T}{B} \right]$,

$$\frac{\text{OPT}_{\text{LP}}^{[T]}(\mathcal{I}_j)}{\mathbb{E}[\text{REW}(\mathcal{A}, \mathcal{I}_j)]} \geq \Omega(\log T). \quad (8.6)$$

Consider two consecutive instances \mathcal{I}_τ and $\mathcal{I}_{\tau+1}$. The outcome matrices in the phases $1, 2, \dots, \tau$ look identical in both these instances. This implies that any randomized algorithm cannot distinguish the two instances (in expectation). Thus, the expected number of times arm A_1 is chosen by algorithm \mathcal{A} in phases $1, 2, \dots, \tau$ is identical. Let α_τ denote the expected number of times \mathcal{A} plays arm A_1 in phase τ . Note that this is the same for all instances $\mathcal{I}_\tau, \mathcal{I}_{\tau+1}, \dots, \mathcal{I}_{T/B}$, as just argued. Thus, we can write

$$\mathbb{E}[\text{REW}(\mathcal{A}, \mathcal{I}_\tau)] = \sum_{j \in [\tau]} j \epsilon \alpha_j. \quad (8.7)$$

Note that the expected number of times arm A_1 is played in phase τ on instances $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{\tau-1}$ does not appear in this expression and thus, is irrelevant for our purposes. Additionally, WLOG we only consider algorithms that exhaust its budget B . Indeed, an algorithm can instead choose only arm A_1 when the number of steps remaining equals its residual budget, without any degradation in the total reward. Combining Equation (8.7) with Lemma 8.9, the LHS in Equation (8.5) is lower-bounded by,

$$\frac{\text{OPT}_{\text{LP}}^{[T]}}{\mathbb{E}[\text{REW}(\mathcal{A})]} \geq \frac{\epsilon B}{2} \cdot \left(\min_{\substack{\alpha \geq 0: \\ \langle \alpha, \mathbf{1} \rangle = B}} \max_{1 \leq \tau \leq T/B} \frac{\tau + 1}{\sum_{j \in [\tau]} j \epsilon \alpha_j} \right). \quad (8.8)$$

We can characterize the optimal solution α in Equation (8.8) as follows. Since the objective is a minimum over $\frac{T}{B}$ convex functions with a single equality constraint on the sum of the variables, from complementary slackness condition the minimum is attained when

$$\text{for each } \tau \in [T/B], \text{ the expression } \left(\sum_{j \in [\tau]} j \alpha_j \right) \cdot \frac{1}{\tau + 1} \text{ is the same.} \quad (8.9)$$

We now prove that Equation (8.9) leads to the following recurrence for the maximizing values of α_j .

$$\forall j \geq 2 \quad \alpha_j = \frac{\alpha_1}{2j}. \quad (8.10)$$

We will prove the recurrence Equation (8.10) via induction. The base case is when $j = 2$. By Equation (8.9),

$$\frac{1}{\alpha_1} = \frac{3/2}{\alpha_1 + 2\alpha_2},$$

which implies that $\alpha_2 = \frac{1}{4}\alpha_1$, and we are done. Now consider the inductive case; let all α up to α_τ satisfy the recurrence Equation (8.10). Consider the instance \mathcal{I}_τ and $\mathcal{I}_{\tau+1}$. From Equation (8.9) we have,

$$\frac{\alpha_1 + \sum_{j=2}^{\tau} \alpha_j/2}{\tau + 1} = \frac{\alpha_1 + \sum_{j=2}^{\tau} \alpha_j/2 + (\tau + 1)\alpha_{\tau+1}}{\tau + 2}.$$

Rearranging, $\alpha_{\tau+1} = \frac{1}{2(\tau+1)}\alpha_1$. This completes the inductive step, and the proof of Equation (8.10).

We complete the proof of the lemma as follows. As argued in Equation (8.10), for the minimum value of $\{\alpha_j\}_{j \in [T/B]}$, the expression $\frac{\epsilon B}{2} \cdot \frac{\tau+1}{\sum_{j \in [\tau]} j \epsilon \alpha_j}$, which is the RHS in Equation (8.8), is the same for all τ and in particular for $\tau = 1$. Substituting $\tau = 1$, this evaluates to B/α_1 . Since $\alpha_1(1 + 1/4 + 1/6 + \dots + B/2T) \leq B$ it follows that $\alpha_1 \leq 2B/H(\frac{T}{B})$, where $H(n)$ denotes the n^{th} Harmonic number. So, the right-hand side of Equation (8.5) is at least $\frac{1}{2}H(\frac{T}{B})$. Finally, $H(n) \geq \ln(n) + \zeta$, where $\zeta = 0.577 \dots$ is the Euler-Masceroni constant. Combining this with Lemma 8.6 we obtain Equation (8.4).

8.3 Best Dynamic Policy: Proof of Theorem 8.1(c) and Theorem 8.3

Consider the following construction of the lower-bound example.

CONSTRUCTION 8.10. *There is one resource with budget B , and two arms, denoted A_0, A_1 . Arm A_0 is the ‘null arm’ that has zero reward and zero consumption. The consumption of arm A_1 is 1 in all rounds. The rewards of A_1 are defined as follows. We partition the time into $\frac{T}{B}$ phases of duration B each (for simplicity, assume that B divides T). We consider $\frac{T}{B}$ problem instances; for each instance \mathcal{I}_τ , $\tau \in [T/B]$ arm A_1 has 0 reward in all phases except phase τ ; in phase τ it has a reward of 1 in each round.*

LEMMA 8.11. *Consider Construction 8.10 with any given time horizon $T \geq 2$ and budget $B \leq \sqrt{T}$. Let ALG be an arbitrary randomized algorithm for BwK. Then for one of the problem instances,*

$$\text{OPT}_{\text{DP}}/\mathbb{E}[\text{REW}] \geq T/B^2. \quad (8.11)$$

PROOF. Let $n = \frac{T}{B}$ be the number of phases in Construction 8.10. Let ALG be a deterministic algorithm. Let REW denote its total reward, and let $\mathbb{E}_\tau[\cdot]$ denote the expectation over the uniform-at-random choice of the problem instance \mathcal{I}_τ . We claim that

$$\text{OPT}_{\text{DP}}/\mathbb{E}_\tau[\text{REW}] \geq T/B^2. \quad (8.12)$$

Assume that ALG maximizes $\mathbb{E}_\tau[\text{REW}]$ (over deterministic algorithms). Then it satisfies the following:

- Within each phase, if ALG ever chooses to play arm A_1 , it does so in the first round of the phase. If it receives a reward of 1 in this round, it plays A_1 for the rest of the phase. Else, it never plays A_1 for the rest of this phase.

For each $\tau \in [n]$, let α_τ denote the number of times ALG chooses arm A_1 in phase τ in problem instance \mathcal{I}_τ . The expected reward of ALG over the uniform-at-random choice of the problem instance \mathcal{I}_τ is $\mathbb{E}[\text{REW}] = \frac{1}{n} \sum_{i \in [n]} \alpha_i$. Let $(\alpha_{\pi(1)}, \alpha_{\pi(2)}, \dots, \alpha_{\pi(k)})$ be the subsequence of $(\alpha_1, \dots, \alpha_n)$ which contains all elements with non-zero values.

The key observation is as follows. The problem instances $\mathcal{I}_{\pi(\tau-1)}$ and $\mathcal{I}_{\pi(\tau)}$ are identical until phase $\pi(\tau-1) - 1$. Since the feedback received by ALG until the first time it chooses arm A_1 in phase $\pi(\tau-1)$ is identical, it follows that $\alpha_{\pi(\tau-1)} - \alpha_{\pi(\tau)} = 1$. Therefore,

$$\sum_{i \in [n]} \alpha_i = \sum_{i \in [k]} \alpha_{\pi(i)} = k \cdot \alpha_{\pi(1)} - \frac{k(k-1)}{2}.$$

Noting that $\alpha_1 \leq B$ and $k \leq \min(B, n) = B$, we have:

$$\mathbb{E}[\text{REW}] \leq \frac{1}{n} \sum_{i \in [n]} \alpha_i < B^2/n = B^3/T.$$

Since $\text{OPT}_{\text{DP}} = B$ for every problem instance \mathcal{I}_τ , Equation (8.12) holds for ALG, and therefore for any other deterministic algorithm. By Yao’s lemma, for every randomized algorithm ALG there exists a problem instance \mathcal{I}_τ such that (8.11) holds. \square

We now use the same construction to prove Theorem 8.3.

PROOF SKETCH OF THEOREM 8.3. We prove the Theorem by contradiction. Let $B \leq \sqrt{T}$. For contradiction, consider an algorithm ALG for cBwK on a policy set Π such that $\text{OPT}_{\text{FD}}(\Pi)/\mathbb{E}[\text{REW}(\text{ALG})] < T/B^2$. We will now use ALG to construct an algorithm \mathcal{A} for the Construction 8.10 such that $\text{OPT}_{\text{DP}}/\mathbb{E}[\text{REW}(\mathcal{A})] < T/B^2$ for every instance. This contradicts Lemma 8.11.

Consider a policy set Π with $|n|$ policies. Every policy $\pi \in \Pi$ maps contexts in the range $[1, T]$ to the action set $\{A_1, A_0\}$. In particular, a policy $\pi_\tau \in \Pi$ maps contexts that lie in the range $[B * (\tau - 1) + 1, B * \tau]$ to arm A_1 and all other contexts to A_0 . \mathcal{A} invokes ALG as a sub-routine with the policy set Π . At each time-step t , \mathcal{A} gives the context $x_t = t$ to ALG and plays the arm chosen by ALG.

For an instance \mathcal{I}_τ in Construction 8.10, $\text{OPT}_{\text{FD}}(\Pi)$ is the total reward obtained by choosing the action given by π_τ in all time-steps. The total reward obtained is B , which equals $\text{OPT}_{\text{DP}}(\mathcal{I}_\tau)$. Therefore, $\text{OPT}_{\text{FD}}(\Pi)/\mathbb{E}[\text{REW}(\text{ALG})] < T/B^2$ implies we have $\text{OPT}_{\text{DP}}/\mathbb{E}[\text{REW}(\mathcal{A})] < T/B^2$ for every instance \mathcal{I}_τ , which is a contradiction. \square

8.4 Best Fixed Arm: Proof of Theorem 8.1(d)

We use the following construction for the lower-bound.

CONSTRUCTION 8.12. *There is one resource with budget B , and K arms denoted by A_1, A_2, \dots, A_K . Arm A_K is the ‘null arm’ that has zero reward and zero consumption. There are K instances in the family. In each instance, the time-steps are divided into T/K equally spaced phases. In instance \mathcal{I}_j , all arms $A_{j'}$ where $j' > j$ have 0 reward and 0 consumption in all time-steps. Consider an instance \mathcal{I}_j for some $j \in [K - 1]$ and an arm $j' \leq j$. Arm $A_{j'}$ has a reward of $\frac{1}{K^{K-j'}}$ and consumption of 1 in all time-steps in phase j' and has a reward of 0 and consumption of 0 in every other time-step. Thus, the rewards and consumption are bounded in the interval $[0, 1]$ for every arm and every time-step in all instances in this family.*

LEMMA 8.13. *Let $T \geq 2$, $2 \leq B \leq T$, $K \geq 3$ be given parameters of the AdversarialBwK problem. We show that there exists a family of instances with $d = 1$ shared resource such that for every randomized algorithm \mathcal{A} we have $\frac{\text{OPT}_{\text{FA}}}{\mathbb{E}[\text{REW}(\mathcal{A})]}$ is at least $\Omega(K)$ on one of these instances.*

PROOF. First note that the best fixed arm in instance \mathcal{I}_j is to pick arm A_j which yields a total reward of $\frac{B}{K^{K-j}}$.

Consider a randomized algorithm \mathcal{A} . Observe that in the first j phases, the instances \mathcal{I}_{j-1} and \mathcal{I}_j have identical outcome matrices. Thus, the expected number of times any arm A_k for $k \in [K]$ is chosen in phases $\{1, 2, \dots, j\}$ should be the same in both the instances. Let α_k denote the expected number of times arm k is played by \mathcal{A} in phase k on instances $\mathcal{I}_k, \mathcal{I}_{k+1}, \dots, \mathcal{I}_{K-1}$.¹⁷ Moreover, we have that $\alpha_1 + \alpha_2, \dots, \alpha_{K-1} \leq B$.

To show the lower-bound, we want to minimize the competitive ratio on every instance for all possible values of $\alpha_1, \alpha_2, \dots, \alpha_{K-1}$. For ease of notation denote $r_j := \frac{1}{K^{K-j}}$. Let $\alpha_{\mathcal{B}}$ denote the set of values to $\{\alpha_k\}_{k \in [K-1]}$ such that $\sum_{k \in [K-1]} \alpha_k \leq B$. Thus,

$$\frac{\text{OPT}_{\text{FA}}}{\mathbb{E}[\text{REW}(\mathcal{A})]} \geq \min_{\alpha_{\mathcal{B}}} \frac{r_k B}{\sum_{j \in [k]} r_j \alpha_j}. \quad (8.13)$$

The ratio is minimized when all ratios in Equation (8.13) are equal. We will show via induction that this yields the following recurrence,

$$\forall k \geq 2 \quad \alpha_k = \left(1 - \frac{r_{k-1}}{r_k}\right) \alpha_1. \quad (8.14)$$

Combining this with the condition that $\sum_{k \in [K-1]} \alpha_k \leq B$, this yields the condition $\alpha_1 \leq \frac{B}{K - \frac{1}{K}}$. Moreover, the minimizing value in Equation (8.13) is $K - \frac{1}{K}$ which proves Lemma 8.13.

¹⁷This has to be the same in all instances since the outcome matrix is identical until phase k in all these instances.

We will now prove the recurrence Equation (8.14). Consider the base case with $k = 2$. Equalizing the first two terms in Equation (8.13) we get

$$\frac{r_1 B}{r_1 \alpha_1} = \frac{r_2 B}{r_1 \alpha_1 + r_2 \alpha_2}.$$

Re-arranging this, we obtain that $\alpha_2 = (1 - \frac{r_1}{r_2})\alpha_1$. We will now prove the inductive case. Let the recurrence be true for all $1 \leq k \leq k'$. Consider the case $k = k' + 1$. Setting the k' and $k' + 1$ ratios in Equation (8.13) equal, we obtain

$$\frac{r_{k'} B}{\sum_{j \in [k']} r_j \alpha_j} = \frac{r_{k'+1} B}{\sum_{j \in [k'+1]} r_j \alpha_j}. \quad (8.15)$$

Moreover, from the inductive hypothesis we have $\alpha_j = (1 - \frac{r_{j-1}}{r_j})\alpha_1$ for every $j \leq k'$. Thus,

$$\begin{aligned} \sum_{j \in [k']} r_j \alpha_j &= r_{k'} \alpha_1 \\ \sum_{j \in [k'+1]} r_j \alpha_j &= r_{k'} \alpha_1 + r_{k'+1} \alpha_{k'+1}. \end{aligned}$$

Plugging this back in Equation (8.15) we get

$$\frac{r_{k'} B}{r_{k'} \alpha_1} = \frac{r_{k'+1} B}{r_{k'} \alpha_1 + r_{k'+1} \alpha_{k'+1}}.$$

Rearranging we get $\alpha_{k'+1} = (1 - \frac{r_{k'}}{r_{k'+1}})\alpha_1$. This completes the induction. \square

9 OPEN QUESTIONS AND FOLLOW-UP WORK

We use essentially the same algorithm, LagrangeBwK, to solve both stochastic and adversarial version of bandits with knapsacks. Yet, we use it with different parameter T_0 (randomly guessed in the adversarial version) and a slightly different definition of the outcome matrices.¹⁸ Can we solve both versions with *exactly* the same algorithm? One concrete goal would be to achieve $O(\log T)$ competitive ratio in the adversarial version, and $o(T)$ regret for the stochastic version in the regime $\min(B, \text{OPT}_{\text{FD}}) \geq \Omega(T)$. A similar ‘‘best of both worlds’’ result has been obtained for bandits without budget/supply constraints: one algorithm that achieves optimal regret rates for both adversarial bandits and stochastic bandits, without knowing which environment it is in [14, 31, 89]. Further developments focused on mostly stochastic environments with a small amount of adversarial behavior [72, 88, 89, 100]; similar questions are relevant to BwK as well.

Given our upper and lower bounds, the competitive ratio $\frac{\text{OPT}_{\text{FD}} - \text{reg}}{\mathbb{E}[\text{REW}]}$ can potentially be improved in several regimes. Some concrete questions left open by our paper are as follows:

- obtain *constant* competitive ratio in the regime $B = \Omega(T)$.
- obtain *sublinear* dependence of the competitive ratio on d , the number of resources.
- obtain *constant* competitive ratio for problem instances with ‘‘large enough’’ OPT_{FD} .
- obtain *optimal* constant competitive ratio when OPT_{FD} is known up to a constant factor.
- obtain competitive ratio in Theorem 5.1 *uniformly* over g_{\min} (i.e., for all g_{\min} simultaneously). Equivalently, obtain competitive ratio $O(\ln(T/\text{OPT}_{\text{FD}}))$.

Several of these questions have been resolved in follow-up work. [66] resolve the optimal dependence on d , achieving competitive ratio $O(\log(d)\log(T))$ via a more careful analysis of

¹⁸Recall that in the stochastic setting there is a ‘dummy resource’ with strictly positive consumption for all arms, whereas in the adversarial version the null arm must have zero consumption for all resources.

LagrangeBwK (among other results), and prove a matching lower bound. Castiglioni et al. [36] use a version of LagrangeBwK to obtain T/B competitive ratio, which is a mere constant when $B = \Omega(T)$. Interestingly, they use fixed parameters $(B_0, T_0) = (B, T)$, without the random guessing in Algorithm 2 or the multi-phased “meta-algorithm” in Algorithm 3; moreover, their result holds with high probability. Castiglioni et al. [36] also analyze the very same algorithm in the stochastic setting, matching our regret bound from Theorem 4.4 and therefore achieving a “best of both worlds” result. Their version of LagrangeBwK optimizes the dual vector $\lambda \in [0, T/B]^d$, whereas ours optimizes λ over all distributions.

Can one still achieve meaningful regret bounds for Adversarial BwK, without the competitive ratio? One way to interpret our impossibility results is that the fixed-distribution benchmark is just too harsh. It could be productive to define a weaker (and perhaps fairer) benchmark for the algorithm to compete against, so as to achieve competitive ratio of 1 relative to this benchmark. Gaitonde et al. [54] achieve one such result for the special case of budget-constrained bidding in a repeated auction.

In terms of extensions to “richer” application scenarios, as in Section 7, Castiglioni et al. [36] spell out two more extensions: to repeated Stackelberg games and to repeated first-price auctions. The main open question is to achieve similar results using a “stochastic” primal algorithm, i.e., a primal algorithm designed (only) for the stochastic version of a particular application scenario.

APPENDICES

A STANDARD TOOLS

Our exposition in the body of the paper relies on some tools that are either known or can easily be derived using standard techniques. We state (and sometimes derive) these tools in this appendix.

A.1 Concentration Inequalities

LEMMA A.1 (AZUMA-HOEFFDING INEQUALITY). *Let Y_1, Y_2, \dots, Y_T be a martingale difference sequence (i.e., $\mathbb{E}[Y_t | Y_1, Y_2, \dots, Y_{t-1}] = 0$). Suppose $|Y_t| \leq c$ for all $t \in \{1, 2, \dots, T\}$. Let $R_{0,\delta}(T) := \sqrt{2Tc^2 \ln(1/\delta)}$. Then for every $\delta > 0$,*

$$\Pr \left[\sum_{t \in [T]} Y_t > R_{0,\delta}(T) \right] \leq \delta.$$

LEMMA A.2 (CHERNOFF-HOEFFDING BOUNDS). *Let X_1, X_2, \dots, X_T be a sequence of independent random variables such that $|X_t| \leq c$ for all $t \in \{1, 2, \dots, T\}$. Let $Z_t := \mathbb{E}[X_t]$. Then for every $\delta > 0$,*

$$\Pr \left[\left| \sum_{t \in [T]} X_t - Z_t \right| > 3 \sqrt{\left(\sum_{t \in [T]} Z_t \right) c^2 \ln(1/\delta)} \right] \leq \delta.$$

A.2 Lagrangians: Proof of Lemma 4.2

Assume one of the resources is the dummy resource, and one of the arms is the null arm. Consider the linear program $\text{LP}_{M,B,T}$, for some outcome matrix M . Let $\mathcal{L} = \mathcal{L}_{M,B,T}$ denote the Lagrange function.

LEMMA A.3 (LEMMA 4.2, RESTATED). *Suppose (X^*, λ^*) is a mixed Nash equilibrium for the Lagrangian game. Then X^* is an optimal solution for the linear program (4.1). Moreover, the minimax value of the Lagrangian game equals the LP value: $\mathcal{L}(X^*, \lambda^*) = \text{OPT}_{\text{LP}}$.*

In what follows we prove Lemma A.3. Writing out the definition of the mixed Nash equilibrium,

$$\mathcal{L}(X^*, \lambda) \geq \mathcal{L}(X^*, \lambda^*) \geq \mathcal{L}(X, \lambda^*) \quad \forall X \in \Delta_K, \lambda \in \Delta_d. \quad (\text{A.1})$$

For brevity, denote $r(X^*) = \sum_{a \in [K]} X^*(a) r(a)$ and $c_i(X^*) = \sum_{a \in [K]} X^*(a) c_i(a)$.

We first state and prove the complementary slackness condition for the Nash equilibrium.

CLAIM A.4. *For every resource $i \in [d]$ we have,*

- (a) $1 - \frac{T}{B} c_i(X^*) \geq 0$,
- (b) $\lambda_i^* > 0 \implies 1 - \frac{T}{B} c_i(X^*) = 0$.

PROOF. **Part (a).** For contradiction, consider resource i that minimizes the left-hand side in (a), and assume that the said left-hand side is strictly negative. We have two cases: either $\lambda_i^* < 1$ or $\lambda_i^* = 1$. When $\lambda_i^* < 1$, consider another distribution $\tilde{\lambda} \in \Delta_d$ such that $\tilde{\lambda}_i = 1$ and $\tilde{\lambda}_{i'} = 0$ for every $i' \neq i$. Note that we have, $\mathcal{L}(X^*, \tilde{\lambda}) < \mathcal{L}(X^*, \lambda^*)$. This contradicts the first inequality in (A.1).

Consider the second case, when $\lambda_i^* = 1$. Then $\mathcal{L}(X^*, \lambda^*) = r(X^*) + 1 - \frac{T}{B} c_i(X^*)$. Consider any arm $a \in [K]$ such that $X^*(a) \neq 0$. Let $\tilde{X} \in \Delta_K$ be another distribution such that $\tilde{X}(a) := 0$ and $\tilde{X}(\text{null}) := X^*(\text{null}) + X^*(a)$ and $\tilde{X}(a') = X^*(a')$ for every $a' \notin \{a, \text{null}\}$. Note that $\tilde{X}(\text{null}) \leq 1$. Since (X^*, λ^*) is a Nash equilibrium, we have that $\mathcal{L}(\tilde{X}, \lambda^*) \leq \mathcal{L}(X^*, \lambda^*)$. This implies that $-X^*(a)r(a) + X^*(a)\frac{T}{B}c_i(a) \leq 0$. Re-arranging we obtain, $\frac{T}{B}c_i(a) \leq r(a) \leq 1$. Thus, we have $1 - \frac{T}{B}c_i(a) \geq 0$.

Since this holds for every $a \in [K]$ with $X^*(a) \neq 0$, we obtain a contradiction:

$$1 - \frac{T}{B} c_i(X^*) = \sum_{a \in [K]} X^*(a) \left(1 - \frac{T}{B} c_i(a) \right) \geq 0.$$

Part (b). For contradiction, assume the statement is false for some resource i . Then, by part (a), $\lambda_i^* > 0$ and $1 - \frac{T}{B}c_i(X^*) > 0$, and consequently $\mathcal{L}(X^*, \lambda^*) > r(X^*)$. Now, consider distribution $\tilde{\lambda}$ which puts probability 1 on the dummy resource. We then have $\mathcal{L}(X^*, \tilde{\lambda}) = r(X^*) < \mathcal{L}(X^*, \lambda^*)$, contradicting the first inequality in Equation (A.1). \square

Let \tilde{X} be some feasible solution for the linear program (4.1). Plugging the feasibility constraints into the definition of the Lagrangian function, $\mathcal{L}(\tilde{X}, \lambda^*) \geq r(\tilde{X})$. Claim A.4(a) implies that X^* is a feasible solution to the linear program (4.1). Claim A.4(b) implies that $\mathcal{L}(X^*, \lambda^*) = r(X^*)$. Thus,

$$r(X^*) = \mathcal{L}(X^*, \lambda^*) \geq \mathcal{L}(\tilde{X}, \lambda^*) \geq r(\tilde{X}).$$

So, X^* is an optimal solution to the LP. In particular, $\text{OPT}_{\text{LP}} = r(X^*) = \mathcal{L}(X^*, \lambda^*)$.

A.3 The Stopped LP for Adversarial BwK: Proof of Equation (5.4)

The proof is similar to prior work [19, 45]. Denote \mathcal{D}_τ to be the set of all distributions over the arms such that for every $\mathbf{p} \in \mathcal{D}_\tau$ we have the following: for every $i \in [d]$ we have $\sum_{t \in [\tau]} \mathbf{p} \cdot \mathbf{c}_{t,i} \leq B$. In other words, \mathcal{D}_τ denotes the set of distributions whose expected stopping time is at least τ . Thus, it immediately follows that $\text{OPT}_{\text{LP}}(\tau) \geq \max_{\mathbf{p} \in \mathcal{D}_\tau} \sum_{t \in [\tau]} \mathbf{p} \cdot \mathbf{r}_t$ since for any given $\mathbf{p} \in \mathcal{D}_\tau$ it is feasible to LP(τ). Thus, $\text{OPT}_{\text{LP}}(\tau)$ is at least the value of any feasible solution $\mathbf{p} \in \mathcal{D}_\tau$. Note that for every fixed distribution $\mathbf{p} \in \Delta_K$, there exists a τ such that either $\mathbf{p} \in \mathcal{D}_\tau$ and $\mathbf{p} \notin \mathcal{D}_{\tau+1}$ or $\mathbf{p} \in \mathcal{D}_T$. Moreover, the total expected reward we can obtain using \mathbf{p} is $\sum_{t \in [\tau]} \mathbf{p} \cdot \mathbf{r}_t$. Thus, $\max_{1 \leq \tau \leq T} \text{OPT}_{\text{LP}}(\tau) \geq \text{OPT}_{\text{FD}}$.

A.4 Regret Minimization in Games: Proof of Lemma 3.1

Let us revisit adversarial online learning, as per Figure 1. Denote the benchmark in Equation (3.2) as

$$\text{OPT}_{\text{AOL}}(T) := \max_{a \in A} \sum_{t \in [T]} f_t(a).$$

Recall that $[b_{\min}, b_{\max}]$ is the payoff range, and denote $\sigma = b_{\max} - b_{\min}$.

LEMMA A.5. *Suppose an algorithm for adversarial online learning satisfies (3.2) for some $\delta > 0$. Then*

$$\Pr \left[\forall \tau \in [T] \text{OPT}_{\text{AOL}}(\tau) - \sum_{t \in [\tau]} f_t \cdot \mathbf{p}_t \leq \sigma \cdot \left(R_{\delta/T}(T) + \sqrt{2T \log(T/\delta)} \right) \right] \geq 1 - 2\delta. \quad (\text{A.2})$$

PROOF. Let us use the stronger regret bound (3.3) implied by (3.2). Note that

$$\mathbb{E}[f_t(a_t) \mid a_1, a_2, \dots, a_{t-1}] = f_t \cdot \mathbf{p}_t.$$

Applying the Azuma-Hoeffding inequality for each $\tau \in [T]$, and taking a union bound, we have

$$\Pr \left[\forall \tau \in [T] \sum_{t \in [\tau]} f_t(a_t) - \sum_{t \in [\tau]} f_t \cdot \mathbf{p}_t \leq \sigma \cdot \sqrt{2T \log(T/\delta)} \right] \geq 1 - \delta. \quad (\text{A.3})$$

Taking a union bound over Equations (A.3) and (3.3) and adding the equations we get Equation (A.2). \square

Remark A.6. For Hedge algorithm, regret bound Equation (A.2) is already proved in [52].

Let $W = \sqrt{2T \log(T/\delta)}$ denote the term from Lemma A.5 in what follows.

We now prove Lemma 3.1, similar to the proof in [51] for the deterministic game. Recall that we take averages up to some fixed round $\tau \in [T]$. We prove that the following two inequalities hold, each with probability at least $1 - \delta$.

$$\frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_{t,1}^T \mathbf{G}_t \mathbf{p}_{t,2} \geq v^* - \sigma \cdot \frac{R_{1, \delta/T}(T) + 2W}{\tau}. \quad (\text{A.4})$$

$$\frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_{t,1}^T \mathbf{G}_t \mathbf{p}_{t,2} \leq \bar{\mathbf{p}}_1^T \mathbf{G} \mathbf{p}_2 + \sigma \cdot \frac{R_{2, \delta/T}(T) + 2W}{\tau} \quad \forall \mathbf{p}_2 \in \Delta_{A_2}. \quad (\text{A.5})$$

Equation (3.5) in Lemma 3.1 follows by adding Equations (A.4) and (A.5).

First we prove Equation (A.4). Following the set of inequalities in Section 2.5 of [51] we have,

$$\begin{aligned} \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_{t,1}^T \mathbf{G}_t \mathbf{p}_{t,2} &\geq_{\text{whp}} \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_1^{*T} \mathbf{G}_t \mathbf{p}_{t,2} - \sigma \cdot \frac{R_{1, \delta/T}(T) + 2W}{\tau} && \text{From Lemma A.5} \\ &\geq_{\text{whp}} \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_1^{*T} \mathbf{G} \mathbf{p}_{t,2} - \sigma \cdot \frac{R_{1, \delta/T}(T) + 2W}{\tau} && \text{From Lemma A.1} \\ &= \max_{\mathbf{p}_1 \in \Delta_{A_1}} \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_1^T \mathbf{G} \mathbf{p}_{t,2} - \sigma \cdot \frac{R_{1, \delta/T}(T) + 2W}{\tau} && \text{From Definition of } \mathbf{p}_1^* \\ &= \max_{\mathbf{p}_1 \in \Delta_{A_1}} \mathbf{p}_1^T \mathbf{G} \bar{\mathbf{p}}_2 - \sigma \cdot \frac{R_{1, \delta/T}(T) + 2W}{\tau} && \text{From Definition of } \bar{\mathbf{p}}_2 \\ &\geq \min_{\mathbf{p}_2 \in \Delta_{A_2}} \max_{\mathbf{p}_1 \in \Delta_{A_1}} \mathbf{p}_1^T \mathbf{G} \mathbf{p}_2 - \sigma \cdot \frac{R_{1, \delta/T}(T) + 2W}{\tau} \end{aligned}$$

Here \leq_{whp} denotes statements that hold with probability at least $1 - \delta$.

Now let us prove (A.5). Fix distribution $\mathbf{p}_2 \in \Delta_{A_2}$. Then:

$$\begin{aligned} \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_{t,1}^\top \mathbf{G}_t \mathbf{p}_{t,2} &\leq_{whp} \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_{t,1}^\top \mathbf{G}_t \mathbf{p}_2 + \sigma \cdot \frac{R_{2,\delta/T}(T) + W}{\tau} && \text{From Lemma A.5} \\ &\leq_{whp} \frac{1}{\tau} \sum_{t \in [\tau]} \mathbf{p}_{t,1}^\top \mathbf{G} \mathbf{p}_2 + \sigma \cdot \frac{R_{2,\delta/T}(T) + 2W}{\tau} && \text{From Lemma A.1} \\ &= \overline{\mathbf{p}}_1^\top \mathbf{G} \mathbf{p}_2 + \sigma \cdot \frac{R_{2,\delta/T}(T) + 2W}{\tau} && \text{From Definition of } \overline{\mathbf{p}}_1. \end{aligned}$$

Taking a union bound over all the four high-probability inequalities, we get the lemma.

B TABLE OF NOTATION

For reference, let us summarize the important notation used across sections.

Notation	Usage
OPT_{FD}	Optimal value of the fixed distribution over arms in hindsight.
OPT_{DP}	Optimal dynamic policy in hindsight.
REW	Total (random) reward obtained by the algorithm.
\mathbf{M}	Outcome matrix; rewards and consumption for every arm.
$\overline{\mathbf{M}}_\tau$	Average of outcome matrices after τ time-steps.
$\overline{\mathbf{M}}_\tau^{\text{ips}}$	Average of outcome matrices estimated using IPS estimates after τ time-steps.
\mathbf{G}	Payoff matrix in the Lagrangian game.
$R_{j,\delta}(\tau)$	Regret of ALG_j with probability at least $1 - \delta$ after τ rounds.
$R_{0,\delta}(\tau)$ or $R_0(\tau)$	Confidence term in the Azuma-Hoeffding inequality.
$U_j(T T_0)$	Regret of ALG_j after T rounds given the parameter T_0 .
$\mathcal{L}(\cdot)$	Lagrange function.
T_0	Parameter in the Lagrangian: T in Stochastic BwK and \widehat{g} in Adversarial BwK.
B_0	Scaled budget: $\frac{B}{\text{ratio}}$ in Adversarial BwK (high-probability), else $B_0 = B$.
$\widehat{g} \in [g_{\min}, g_{\max}]$	Guess and range thereof in Adversarial BwK.
κ	Multiplicative factor with which guess is increased.
$\text{OPT}_{\text{LP}}^{[\tau]}$	Best objective of the τ stopped LPs (i.e., stopped at times $1, 2, \dots, \tau$).
$\text{LP}_{\overline{\mathbf{M}}_\tau, B, \tau}$	Linear program corresponding to the average outcome matrix $\overline{\mathbf{M}}_\tau$.
$\text{OPT}_{\text{LP}}(\overline{\mathbf{M}}_\tau, B, \tau)$	Optimal value of $\text{LP}_{\overline{\mathbf{M}}_\tau, B, \tau}$.

ACKNOWLEDGMENTS

We are grateful to Sahil Singla and Thomas Kesselheim for pointing out the reduction in Remark 5.6, and an inefficiency in our original analysis in Section 5.1. We are also grateful to Omid Sadeghi and the JACM reviewers for pointing out several typos and inaccuracies. We thank Robert Kleinberg, Akshay Krishnamurthy, Steven Wu, and Chicheng Zhang for many insightful conversations on online machine learning.

REFERENCES

- [1] Jacob D. Abernethy and Jun-Kun Wang. 2017. On Frank-Wolfe and equilibrium computation. In *Advances in Neural Information Processing Systems (NIPS)*. 6584–6593.
- [2] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. 2017. A reductions approach to fair classification. *Fairness, Accountability, and Transparency in Machine Learning (FATML)* (2017).
- [3] Alekh Agarwal, Miroslav Dudík, Satyen Kale, John Langford, and Robert E. Schapire. 2012. Contextual bandit learning with predictable rewards. In *15th Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*. 19–26.

- [4] Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. 2014. Taming the monster: A fast and simple algorithm for contextual bandits. In *31st Intl. Conf. on Machine Learning (ICML)*.
- [5] Shipra Agrawal and Nikhil R. Devanur. 2014. Bandits with concave rewards and convex knapsacks. In *15th ACM Conf. on Economics and Computation (ACM-EC)*.
- [6] Shipra Agrawal and Nikhil R. Devanur. 2016. Linear contextual bandits with knapsacks. In *29th Advances in Neural Information Processing Systems (NIPS)*.
- [7] Shipra Agrawal, Nikhil R. Devanur, and Lihong Li. 2016. An efficient algorithm for contextual bandits with knapsacks, and an extension to concave objectives. In *29th Conf. on Learning Theory (COLT)*.
- [8] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. 2014. A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62, 4 (2014), 876–890.
- [9] Noga Alon, Baruch Awerbuch, and Yossi Azar. 2003. The online set cover problem. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*. ACM, 100–105.
- [10] Sanjeev Arora, Elad Hazan, and Satyen Kale. 2012. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing* 8, 1 (2012), 121–164.
- [11] Sanjeev Arora, Satish Rao, and Umesh Vazirani. 2009. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)* 56, 2 (2009), 5.
- [12] Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. 2011. Minimax policies for combinatorial prediction games. In *24th Conf. on Learning Theory (COLT)*. 107–132.
- [13] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. 2002. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.* 32, 1 (2002), 48–77. Preliminary version in *36th IEEE FOCS*, 1995.
- [14] Peter Auer and Chao-Kai Chiang. 2016. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *29th Conf. on Learning Theory (COLT)*.
- [15] Baruch Awerbuch and Yossi Azar. 1997. Buy-at-bulk network design. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*. IEEE, 542–547.
- [16] Yossi Azar, Niv Buchbinder, T. H. Hubert Chan, Shahar Chen, Ilan Reuven Cohen, Anupam Gupta, Zhiyi Huang, Ning Kang, Viswanath Nagarajan, and Joseph Naor. 2016. Online algorithms for covering and packing problems with convex objectives. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 148–157.
- [17] Moshe Babaioff, Shaddin Dughmi, Robert D. Kleinberg, and Aleksandrs Slivkins. 2015. Dynamic pricing with limited supply. *ACM Trans. on Economics and Computation* 3, 1 (2015), 4. Special issue for *13th ACM EC*, 2012.
- [18] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. 2012. Learning on a budget: Posted price mechanisms for online procurement. In *13th ACM Conf. on Electronic Commerce (ACM-EC)*. 128–145.
- [19] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. 2018. Bandits with knapsacks. *J. of the ACM* 65, 3 (2018), 13:1–13:55. Preliminary version in *FOCS 2013*.
- [20] Ashwinkumar Badanidiyuru, John Langford, and Aleksandrs Slivkins. 2014. Resourceful contextual bandits. In *27th Conf. on Learning Theory (COLT)*.
- [21] Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. 2011. A polylogarithmic-competitive algorithm for the k -server problem. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. IEEE, 267–276.
- [22] Ahron Ben-Tal and Arkadi Nemirovski. 2001. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. Vol. 2. Siam.
- [23] Dirk Bergemann and Juuso Välimäki. 2006. Bandit problems. In *The New Palgrave Dictionary of Economics, 2nd ed.*, Steven Durlauf and Larry Blume (Eds.). Macmillan Press.
- [24] Omar Besbes and Assaf Zeevi. 2009. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research* 57, 6 (2009), 1407–1420.
- [25] Omar Besbes and Assaf J. Zeevi. 2012. Blind network revenue management. *Operations Research* 60, 6 (2012), 1537–1550.
- [26] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert E. Schapire. 2011. Contextual bandit algorithms with supervised learning guarantees. In *14th Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*.
- [27] Stephen Boyd and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.
- [28] Sébastien Bubeck and Nicolò Cesa-Bianchi. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5, 1 (2012), 1–122. Published with *Now Publishers* (Boston, MA, USA). Also available at <https://arxiv.org/abs/1204.5721>.
- [29] Sébastien Bubeck, Ofer Dekel, Tomer Koren, and Yuval Peres. 2015. Bandit convex optimization: \sqrt{T} regret in one dimension. In *28th Conf. on Learning Theory (COLT)*. 266–278.
- [30] Sébastien Bubeck, Yin Tat Lee, and Ronen Eldan. 2017. Kernel-based methods for bandit convex optimization. In *49th ACM Symp. on Theory of Computing (STOC)*. ACM, 72–85.

- [31] Sébastien Bubeck and Aleksandrs Slivkins. 2012. The best of both worlds: Stochastic and adversarial bandits. In *25th Conf. on Learning Theory (COLT)*.
- [32] Niv Buchbinder and Joseph Seffi Naor. 2009. The design of competitive online algorithms via a primal–dual approach. *Foundations and Trends® in Theoretical Computer Science* 3, 2–3 (2009), 93–263.
- [33] Niv Buchbinder and Joseph (Seffi) Naor. 2009. Online primal-dual algorithms for covering and packing. *Math. Oper. Res.* 34, 2 (May 2009), 270–286.
- [34] Adrian Rivera Cardoso, Jacob D. Abernethy, He Wang, and Huan Xu. 2019. Competing against Nash equilibria in adversarially changing zero-sum games. In *36th Intl. Conf. on Machine Learning (ICML)*. 921–930.
- [35] Adrian Rivera Cardoso, He Wang, and Huan Xu. 2018. Online saddle point problem with applications to constrained online convex optimization. *arXiv preprint arXiv:1806.08301* (2018).
- [36] Matteo Castiglioni, Andrea Celli, and Christian Kroer. 2022. Online learning with knapsacks: The best of both worlds. In *39th Intl. Conf. on Machine Learning (ICML)*.
- [37] Nicolò Cesa-Bianchi and Gábor Lugosi. 2006. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, UK.
- [38] Moses Charikar and Balaji Raghavachari. 1998. The finite capacity dial-a-ride problem. In *Proceedings 39th Annual Symposium on Foundations of Computer Science*. IEEE, 458–467.
- [39] Tianyi Chen and Georgios B. Giannakis. 2018. Bandit convex optimization for scalable and dynamic IoT management. *IEEE Internet of Things Journal* (2018).
- [40] Tianyi Chen, Qing Ling, and Georgios B. Giannakis. 2017. An online convex optimization approach to proactive network resource allocation. *IEEE Transactions on Signal Processing* 65, 24 (2017), 6350–6364.
- [41] Paul Christiano, Jonathan A. Kelner, Aleksander Madry, Daniel A. Spielman, and Shang-Hua Teng. 2011. Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs. In *43rd ACM Symp. on Theory of Computing (STOC)*. ACM, 273–282.
- [42] Richard Combes, Chong Jiang, and Rayadurgam Srikant. 2015. Bandits with budgets: Regret lower bounds and optimal algorithms. *ACM SIGMETRICS Performance Evaluation Review* 43, 1 (2015), 245–257.
- [43] Constantinos Daskalakis, Alan Deckelbaum, and Anthony Kim. 2015. Near-optimal no-regret algorithms for zero-sum games. *Games and Economic Behavior* 92 (2015), 327–348. Preliminary version in *ACM-SIAM SODA 2011*.
- [44] Nikhil R. Devanur and Thomas P. Hayes. 2009. The AdWords problem: Online keyword matching with budgeted bidders under random permutations. In *10th ACM Conf. on Electronic Commerce (ACM-EC)*. 71–78.
- [45] Nikhil R. Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A. Wilkens. 2011. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *12th ACM Conf. on Electronic Commerce (ACM-EC)*. 29–38.
- [46] Wenkui Ding, Tao Qin, Xu-Dong Zhang, and Tie-Yan Liu. 2013. Multi-armed bandit with budget constraint and variable costs. In *27th AAAI Conference on Artificial Intelligence (AAAI)*.
- [47] Miroslav Dudík, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and Tong Zhang. 2011. Efficient optimal learning for contextual bandits. In *27th Conf. on Uncertainty in Artificial Intelligence (UAI)*.
- [48] Jon Feldman, Monika Henzinger, Nitish Korula, Vahab S. Mirrokni, and Clifford Stein. 2010. Online stochastic packing applied to display ad allocation. In *18th Annual European Symp. on Algorithms (ESA)*. 182–194.
- [49] Amos Fiat, Richard M. Karp, Michael Luby, Lyle A. McGeoch, Daniel D. Sleator, and Neal E. Young. 1991. Competitive paging algorithms. *Journal of Algorithms* 12, 4 (1991), 685–699.
- [50] Abraham Flaxman, Adam Kalai, and H. Brendan McMahan. 2005. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *16th ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 385–394.
- [51] Yoav Freund and Robert E. Schapire. 1996. Game theory, on-line prediction and boosting. In *9th Conf. on Learning Theory (COLT)*. 325–332.
- [52] Yoav Freund and Robert E. Schapire. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* 55, 1 (1997), 119–139.
- [53] Yoav Freund and Robert E. Schapire. 1999. Adaptive game playing using multiplicative weights. *Games and Economic Behavior* 29, 1–2 (1999), 79–103.
- [54] Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. 2022. Budget Pacing in Repeated Auctions: Regret and Efficiency without Convergence. (2022). Working paper, available at <https://arxiv.org/abs/2205.08674>.
- [55] John Gittins, Kevin Glazebrook, and Richard Weber. 2011. *Multi-Armed Bandit Allocation Indices* (2nd ed.). John Wiley & Sons, Hoboken, NJ, USA.
- [56] Sudipta Guha and Kamesh Munagala. 2007. Multi-armed bandits with metric switching costs. In *36th Intl. Colloquium on Automata, Languages and Programming (ICALP)*. 496–507.
- [57] Anupam Gupta, Ravishankar Krishnaswamy, Marco Molinaro, and R. Ravi. 2011. Approximation algorithms for correlated knapsacks and non-martingale bandits. In *52nd IEEE Symp. on Foundations of Computer Science (FOCS)*. 827–836.

- [58] András György, Levente Kocsis, Ivett Szabó, and Csaba Szepesvári. 2007. Continuous time associative bandit problems. In *20th Intl. Joint Conf. on Artificial Intelligence (IJCAI)*. 830–835.
- [59] András György, Tamás Linder, Gábor Lugosi, and György Ottucsák. 2007. The on-line shortest path problem under partial monitoring. *J. of Machine Learning Research (JMLR)* 8 (2007), 2369–2403.
- [60] Elad Hazan. 2015. Introduction to online convex optimization. *Foundations and Trends in Optimization* 2, 3–4 (2015), 157–325. Published with Now Publishers (Boston, MA, USA). Also available at <https://arxiv.org/abs/1909.05207>.
- [61] Elad Hazan and Kfir Y. Levy. 2014. Bandit convex optimization: Towards tight bounds. In *27th Advances in Neural Information Processing Systems (NIPS)*. 784–792.
- [62] Justin Hsu, Zhiyi Huang, Aaron Roth, and Zhiwei Steven Wu. 2016. Jointly private convex programming. In *27th ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 580–599.
- [63] David S. Johnson. 1974. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences* 9, 3 (1974), 256–278.
- [64] Satyen Kale, Lev Reyzin, and Robert E. Schapire. 2010. Non-stochastic bandit slate problems. In *24th Advances in Neural Information Processing Systems (NIPS)*. 1054–1062.
- [65] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *35th Intl. Conf. on Machine Learning (ICML)*. 2564–2572.
- [66] Thomas Kesselheim and Sahil Singla. 2020. Online learning with vector costs and bandits with knapsacks. In *33rd Conf. on Learning Theory (COLT)*. 2286–2305.
- [67] Robert Kleinberg. 2004. Nearly tight bounds for the continuum-armed bandit problem. In *18th Advances in Neural Information Processing Systems (NIPS)*.
- [68] John Langford and Tong Zhang. 2007. The epoch-greedy algorithm for contextual multi-armed bandits. In *21st Advances in Neural Information Processing Systems (NIPS)*.
- [69] Tor Lattimore and Csaba Szepesvári. 2020. *Bandit Algorithms*. Cambridge University Press, Cambridge, UK. Versions available at <https://banditalgs.com/>. since 2018.
- [70] Nick Littlestone and Manfred K. Warmuth. 1994. The weighted majority algorithm. *Information and Computation* 108, 2 (1994), 212–260.
- [71] László Lovász. 1975. On the ratio of optimal integral and fractional covers. *Discrete Mathematics* 13, 4 (1975), 383–390.
- [72] Thodoris Lykouris, Vahab Mirrokni, and Renato Paes-Leme. 2018. Stochastic bandits robust to adversarial corruptions. In *50th ACM Symp. on Theory of Computing (STOC)*.
- [73] Mehrdad Mahdavi, Rong Jin, and Tianbao Yang. 2012. Trading regret for efficiency: Online convex optimization with long term constraints. *J. of Machine Learning Research (JMLR)* 13, Sep. (2012), 2503–2528.
- [74] Mehrdad Mahdavi, Tianbao Yang, and Rong Jin. 2013. Stochastic convex optimization with multiple objectives. In *Advances in Neural Information Processing Systems (NIPS)*. 1115–1123.
- [75] Aranyak Mehta. 2013. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science* 8, 4 (2013), 265–368.
- [76] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. AdWords and generalized online matching. *J. ACM* 54, 5 (2007), 22.
- [77] Marco Molinaro and R. Ravi. 2012. Geometry of online packing linear programs. In *39th Intl. Colloquium on Automata, Languages and Programming (ICALP)*. 701–713.
- [78] Michael J. Neely and Hao Yu. 2017. Online convex optimization with time-varying constraints. *arXiv preprint arXiv:1702.04783* (2017).
- [79] Gergely Neu and Gábor Bartók. 2016. Importance weighting without importance weights: An efficient algorithm for combinatorial semi-bandits. *J. of Machine Learning Research (JMLR)* 17, 1 (2016), 5355–5375.
- [80] Alexander Rakhlin and Karthik Sridharan. 2013. Optimization, learning, and games with predictable sequences. In *27th Advances in Neural Information Processing Systems (NIPS)*. 3066–3074.
- [81] Alexander Rakhlin and Karthik Sridharan. 2016. BISTRO: An efficient relaxation-based method for contextual bandits. In *33rd Intl. Conf. on Machine Learning (ICML)*.
- [82] Anshuka Rangi, Massimo Franceschetti, and Long Tran-Thanh. 2019. Unifying the stochastic and the adversarial bandits with knapsack. In *28th Intl. Joint Conf. on Artificial Intelligence (IJCAI)*. 3311–3317.
- [83] Ryan Rogers, Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. 2015. Inducing approximately optimal flow using truthful mediators. In *16th ACM Conf. on Electronic Commerce (ACM-EC)*. 471–488.
- [84] Aaron Roth, Aleksandrs Slivkins, Jonathan Ullman, and Zhiwei Steven Wu. 2017. Multidimensional dynamic pricing for welfare maximization. In *18th ACM Conf. on Electronic Commerce (ACM-EC)*. 519–536.
- [85] Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. 2016. Watch and learn: Optimizing from revealed preferences feedback. In *48th ACM Symp. on Theory of Computing (STOC)*. 949–962.
- [86] Karthik Abinav Sankararaman and Aleksandrs Slivkins. 2018. Combinatorial semi-bandits with knapsacks. In *Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*. 1760–1770.

- [87] Robert E. Schapire and Yoav Freund. 2012. *Boosting: Foundations and Algorithms*. The MIT Press.
- [88] Yevgeny Seldin and Gábor Lugosi. 2017. An improved parametrization and analysis of the EXP3++ algorithm for stochastic and adversarial bandits. In *30th Conf. on Learning Theory (COLT)*.
- [89] Yevgeny Seldin and Aleksandrs Slivkins. 2014. One practical algorithm for both stochastic and adversarial bandits. In *31st Intl. Conf. on Machine Learning (ICML)*.
- [90] Adish Singla and Andreas Krause. 2013. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *22nd Intl. World Wide Web Conf. (WWW)*. 1167–1178.
- [91] Aleksandrs Slivkins. 2013. Dynamic Ad Allocation: Bandits with Budgets. A technical report on arxiv.org/abs/1306.0155. (June 2013).
- [92] Aleksandrs Slivkins. 2019. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning* 12, 1–2 (Nov. 2019), 1–286. Published with Now Publishers (Boston, MA, USA). Also available at <https://arxiv.org/abs/1904.07272>. Latest online revision: Jan. 2022.
- [93] Vasilis Syrgkanis, Akshay Krishnamurthy, and Robert E. Schapire. 2016. Efficient algorithms for adversarial contextual learning. In *33rd Intl. Conf. on Machine Learning (ICML)*.
- [94] Vasilis Syrgkanis, Haipeng Luo, Akshay Krishnamurthy, and Robert E. Schapire. 2016. Improved regret bounds for oracle-based adversarial contextual bandits. In *29th Advances in Neural Information Processing Systems (NIPS)*.
- [95] Long Tran-Thanh, Archie Chapman, Enrique Munoz de Cote, Alex Rogers, and Nicholas R. Jennings. 2010. ϵ -first policies for budget-limited multi-armed bandits. In *24th AAAI Conference on Artificial Intelligence (AAAI)*. 1211–1216.
- [96] Long Tran-Thanh, Archie Chapman, Alex Rogers, and Nicholas R. Jennings. 2012. Knapsack based optimal policies for budget-limited multi-armed bandits. In *26th AAAI Conference on Artificial Intelligence (AAAI)*. 1134–1140.
- [97] Seun Umboh. 2015. Online network design algorithms via hierarchical decompositions. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 1373–1387.
- [98] Jun-Kun Wang and Jacob D. Abernethy. 2018. Acceleration through optimistic no-regret dynamics. In *31st Advances in Neural Information Processing Systems (NIPS)*. 3828–3838.
- [99] Zizhuo Wang, Shiming Deng, and Yinyu Ye. 2014. Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Operations Research* 62, 2 (2014), 318–331.
- [100] Chen-Yu Wei and Haipeng Luo. 2018. More adaptive algorithms for adversarial bandits. In *31st Conf. on Learning Theory (COLT)*.
- [101] David P. Williamson and David B. Shmoys. 2011. *The Design of Approximation Algorithms*. Cambridge University Press.

Received 12 November 2019; revised 23 September 2021; accepted 18 July 2022