## Bandits with Knapsacks beyond the Worst Case (Supplementary Materials)

## Contents

A Motivating examples with $d=2$ and small number of arms ..... 16
B Confidence bounds in UcbBwK ..... 16
C LP Sensitivity: proof of Lemma 3.3 ..... 17
D Various technicalities from Sections 3 ..... 19
D. 1 Standard tools ..... 19
D. 2 Proof of Eq. (3.6) ..... 19
D. 3 Proof of Eq. (3.9) ..... 19
D. 4 Lower bound on Lagrange gap: Proof of Eq. (3.10) ..... 20
D. 5 Martingale arguments: Proof of Eq. (3.8) ..... 20
E Proof of Theorem 3.10: generic $\sqrt{T}$ lower bound ..... 21
F Proof of Theorem 3.9(b): $\sqrt{T}$ lower bound for $d>2$ ..... 24
G Simple regret: proof of Theorem 4.1 ..... 25
G. 1 Confidence sums ..... 25
G. 2 Connecting LP-gap and the confidence radius ..... 26
G. 3 Finishing the proof of Theorem 4.1 ..... 27
G. 4 The standard confidence-sum bound: proof of Eq. (G.1) ..... 27
H Reduction from BwK to bandits ..... 27
H. 1 Linear Contextual Bandits with Knapsacks (LinCBwK) ..... 28
H. 2 Combinatorial Semi-bandits with Knapsacks (SemiBwK) ..... 28
H. 3 Multinomial-logit Bandits with Knapsacks (MnlBwK) ..... 29
H. 4 Computational issues ..... 30

## A Motivating examples with $d=2$ and small number of arms

We provide direct motivation for Theorem 3.2, our positive result for $O(\log T)$ regret. Recall that Theorem 3.2 only holds with $d=2$ resources, and is only meaningful with a reasonably small number of arms $K$ (because the regret bounds are linear in $K$ ). Such problems arise in many motivating applications of BwK, e.g., as listed in [14, 16]. Below we spell out several stylized examples.
In dynamic assortment [50, 7, 26], an algorithm is a seller which chooses among possible assortments of products. In each round, a customer arrives, the algorithm chooses an assortment, and offers this assortment for sale at an exogenously fixed price. If a sale happens, the algorithm receives revenue and consumes some amount of inventory. The following version features $d=2$ and non-huge $K$ : there are $K$ possible offerings for sale, and a limited amount of "raw material" used to manufacture them. Each offering, if sold, consumes some pre-fixed amount of this raw material. ${ }^{8}$
The "inverted" dynamic assortment problem takes the procurement perspective. An algorithm is a budget-limited contractor which chooses among $K$ possible types of offers, e.g., different items to procure from vendors, or different tasks to complete in an online labor market. In each round, a new agent arrives, the algorithm chooses an offer and presents it to the customer at an exogenously fixed price. If the offer is accepted, the contractor receives some utility (i.e., reward) and spends the corresponding amount of money.

In dynamic pricing $[19,12,20,59]$ an algorithm is a seller with limited supply of some product, and chooses a price in each round. If this price is accepted, a sale happens, and algorithm receives revenue and spends inventory. Of our interest is the case when the set of possible prices is small and exogenously fixed, e.g., there are a few possible discount levels. Likewise, in dynamic procurement [13,52, 14], an algorithm is a budget-limited contractor who continuously procures some product or service. The algorithm chooses a price in each round. If this price is accepted, a transaction happens, so that the algorithm receives an "item" (i.e., reward of 1 ) and spends the corresponding amount of money. We focus on the case when there are only a few possible prices, e.g., exogenously fixed levels of premium or surcharge.

Our last example concerns fault-tolerance in systems. Consider a system, either physical or computational, which experiments with different possible policies to process incoming requests. In each time step, it chooses one of the possible policies, and observes the outcome (and there are no lingering effects, e.g., no persistent "system state" that changes over time). The outcome consists of utility for performance-as-usual (i.e., reward), and penalty for various mistakes or faults. Fault-tolerance requirement is expressed as a a "budget" on the total penalty accrued by the algorithm.

## B Confidence bounds in UcbBwK

Let us fill in the exact specification of the confidence bounds in the UcbBwK algorithm. (This is for the sake of completeness only; as pointed out in Preliminaries, these details do not affect our analysis.)

Confidence radius. Given an unknown quantity $\mu$ and its estimator $\hat{\mu}$, a confidence radius is an observable high-confidence upper bound on $|\mu-\widehat{\mu}|$. More formally, it is some quantity $\operatorname{Rad} \in \mathbb{R}_{\geq 0}$ such that it is computable from the algorithm's observations, and $|\mu-\widehat{\mu}| \leq \operatorname{Rad}$ with probability (say) at least $1-1 / T^{3}$. Throughout, the estimator $\widehat{\mu}$ is a sample average over all available observations pertaining to $\mu$, unless specified otherwise.

Following the prior work on $\operatorname{BwK}[12,16,3]$, we use the confidence radius from [38]:

$$
\begin{equation*}
f_{\mathrm{rad}}(\widehat{\mu}, N):=\min \left(1, \sqrt{\frac{C_{\mathrm{rad}} \widehat{\mu}}{\max (1, N)}}+\frac{C_{\mathrm{rad}}}{\max (1, N)}\right), \text { where } C_{\mathrm{rad}}=3 \cdot \log (K d T), \tag{B.1}
\end{equation*}
$$

and $N$ is the number of samples. If $\widehat{\mu}$ is a sample average of $N$ independent random variables with support in $[0,1]$, and $\mu=\mathbb{E}[\mu]$, then with probability at least $1-(K d t)^{-2}$ we have

$$
\begin{equation*}
|\widehat{\mu}-\mu| \leq f_{\mathrm{rad}}(\widehat{\mu}, N) \leq 3 f_{\mathrm{rad}}(\mu, N) \tag{B.2}
\end{equation*}
$$

For each arm, we use this confidence radius separately for expected reward of this arm, and expected consumption of each resource. $x$

[^0]Confidence bounds. Fix arm $a \neq$ null, round $t$, and resource $j \neq$ time.
Let $S_{t}(a)=\left\{s<t: a_{s}=a\right\}$ be the set of all previous rounds in which this arm has been chosen, and let $N_{t}(a)=\left|S_{t}(a)\right|$. Let

$$
\begin{equation*}
\hat{r}_{t}(a):=\frac{1}{t} \sum_{s \in S_{t}(a)} r_{s}(a) \quad \text { and } \quad \hat{c}_{j, t}(a):=\frac{1}{t} \sum_{s \in S_{t}(a)} c_{j, s}(a) \tag{B.3}
\end{equation*}
$$

denote, resp., the sample average of reward and resource- $j$ consumption of this arm so far.
Define the confidence radii $\operatorname{Rad}_{0, t}(a)$ and $\operatorname{Rad}_{j, t}(a)$ for, resp., expected reward $r(a)$ and resource consumption $c_{j}(a)$, and the associated upper/lower confidence bounds:

$$
\begin{align*}
r_{t}^{ \pm}(a) & =\operatorname{proj}\left(\hat{r}_{t}(a) \pm \operatorname{Rad}_{0, t}(a)\right), \\
c_{j, t}^{ \pm}(a) & =\operatorname{proj}\left(\hat{c}_{j, t}(a) \pm \operatorname{Rad}_{j, t}(a)\right), \tag{B.4}
\end{align*} \quad \operatorname{Rad}_{0, t}(a):=f_{\text {rad }}\left(\hat{r}_{t}(a), N_{t}(a)\right), ~=\operatorname{Rad}_{j, t}(a):=f_{\mathrm{rad}}\left(\hat{c}_{j, t}(a), N_{t}(a)\right),
$$

where $\operatorname{proj}(x):=\arg \min _{y \in[0,1]}|y-x|$ denotes the projection into $[0,1]$. Then, the event

$$
\begin{equation*}
r(a) \in\left[r_{t}^{-}(a), r_{t}^{+}(a)\right] \text { and } c_{j}(a) \in\left[c_{j, t}^{-}(a), c_{j, t}^{+}(a)\right], \quad \forall a \in[K], j \in[d-1] \tag{B.5}
\end{equation*}
$$

holds for each round $t$ with probability (say) at least $1-\frac{\log (K d T)}{T^{4}}$ [12].
Note that all confidence radii in (B.4) are upper-bounded by

$$
\begin{equation*}
\operatorname{Rad}_{t}(a):=f_{\mathrm{rad}}\left(1, N_{t}(a)\right), \tag{B.6}
\end{equation*}
$$

which is a version of a more standard confidence radius $\widetilde{O}\left(1 / \sqrt{N_{t}(a)}\right)$.
There is no uncertainty on the time resource and the null arm. So, we set $\operatorname{Rad}_{\text {time }, t}(\cdot)=0$ and $c_{\text {time }, t}^{ \pm}(\cdot)=B / T$, and $\operatorname{Rad}_{0, t}($ null $)=\operatorname{Rad}_{j, t}($ null $)=r^{ \pm}($null $)=c_{j, t}^{ \pm}($null $)=0$.

## C LP Sensitivity: proof of Lemma 3.3

We focus on the sensitivity of the support of the optimal solution. We build on some well-known results, which we state below in a convenient form (and provide a proof for completeness). We use the textbook material from Bertsimas and Tsitsiklis [18].

Throughout this appendix, we consider a best-arm-optimal problem instance with best arm $a^{*}$. Let $\boldsymbol{X}^{*}$ denote the optimal solution for the linear program (2.2). Recall that the support of $X^{*}$ is either $\left\{a^{*}\right\}$ or $\left\{a^{*}\right.$, null $\}$. We consider perturbations in the rescaled $L P$ :

$$
\begin{array}{ll}
\text { maximize } & \boldsymbol{X} \cdot \boldsymbol{r} \\
& \boldsymbol{X} \in[0,1]^{K} \\
\forall j \in[d-1] & \boldsymbol{X} \cdot \mathbf{1}=1  \tag{C.1}\\
& \boldsymbol{X} \cdot \boldsymbol{c}_{j} \leq(B / T)\left(1-\eta_{\mathrm{LP}}\right) \\
& \boldsymbol{X} \cdot \boldsymbol{c}_{d} \leq B / T
\end{array}
$$

Recall that $\boldsymbol{r}, \boldsymbol{c}_{j} \in[0,1]^{K}$ are vectors of expected rewards and expected consumption of resource $j$. The $d$-th resource is time. The rescaling parameter $\eta_{\mathrm{LP}}$ is given in Eq. (2.5).

Let $0 \mathrm{PT}_{\mathrm{LP}}^{\mathrm{sc}}$ denote the value of this LP ; it is easy to see that $\mathrm{OPT}_{\mathrm{LP}}^{\mathrm{sc}}=\left(1-\eta_{\mathrm{LP}}\right) \mathrm{OPT}_{\mathrm{LP}}$.
We observe that $a^{*}$ is the best arm for the rescaled LP, too, because $G_{\text {LAG }}$ is large enough. Call a distribution over arms null-degenerate if its support includes exactly one non-null arm.
Claim C.1. The rescaled $L P(C .1)$ has a null-degenerate optimal solution with non-null arm $a^{*}$.

Proof. From the theory in [18, Ch.5], if the optimal basis to LP (2.2) remains feasible to the rescaled LP (C.1) then the basis is also optimal to this LP. This is because LP (C.1) is obtained by a small perturbation to the right-hand side values in LP (2.2). Let $\boldsymbol{X}^{*}$ denote the optimal solution to LP (2.2). From assumption this is a null-degenerate optimal solution. Using the same analysis in [18, Ch. 4.4] we only have to show that the perturbation is smaller than $X^{*}\left(a^{*}\right)$. Since the perturbation is $\frac{B \eta_{\mathrm{LP}}}{T} \leq \frac{3 \sqrt{B} \log (K T d)}{T}$ while $X^{*}\left(a^{*}\right)>\frac{3 \sqrt{B} \log (K T d)}{T}$, this perturbation does not change the basis. Thus, the rescaled LP has a null-degenerate optimal solution.

Claim C.2. Let $\boldsymbol{\lambda}^{*}$ denote the vector of the optimal dual solution to the $L P(2.2)$. Then

$$
\begin{equation*}
G_{L A G}(a)=\frac{T}{B} \sum_{j \in[d]} \lambda_{j}^{*} c_{j}(a)-r(a) . \tag{C.2}
\end{equation*}
$$

Proof. From Eq. (3.1) we have the following.

$$
\begin{aligned}
G_{\mathrm{LAG}}(a) & :=\mathcal{L}\left(\boldsymbol{X}^{*}, \boldsymbol{\lambda}^{*}\right)-\mathcal{L}\left(\boldsymbol{X}_{a}, \boldsymbol{\lambda}^{*}\right) \\
& =\boldsymbol{r}\left(\boldsymbol{X}^{*}\right)-\frac{T}{B} \sum_{j \in[d]} \lambda_{j}^{*} \boldsymbol{c}_{j}\left(\boldsymbol{X}^{*}\right)+\frac{T}{B} \sum_{j \in[d]} \lambda_{j}^{*} c_{j}(a)-r(a)
\end{aligned}
$$

Consider the dual of the LP (2.2). It can be seen that the objective of this dual is $\sum_{j \in[d]} \lambda_{j}$. It follows that $\mathrm{OPT}_{\mathrm{LP}}=\sum_{j \in[d]} \lambda_{j}^{*}$ by strong duality [22, Section 5.2.3]. As proved in [35], $\mathcal{L}\left(\boldsymbol{X}^{*}, \boldsymbol{\lambda}^{*}\right)=\mathrm{OPT}_{\mathrm{LP}}$. Thus,

$$
\sum_{j \in[d]} \lambda_{j}^{*}=\mathrm{OPT}_{\mathrm{LP}}=\mathcal{L}\left(\boldsymbol{X}^{*}, \boldsymbol{\lambda}^{*}\right)=\boldsymbol{r}\left(\boldsymbol{X}^{*}\right)-\frac{T}{B} \sum_{j \in[d]} \lambda_{j}^{*} \boldsymbol{c}_{j}\left(\boldsymbol{X}^{*}\right)+\sum_{j \in[d]} \lambda_{j}^{*}
$$

Therefore, $\boldsymbol{r}\left(\boldsymbol{X}^{*}\right)=\frac{T}{B} \sum_{j \in[d]} \lambda_{j}^{*} \boldsymbol{c}_{j}\left(\boldsymbol{X}^{*}\right)$, which implies (C.2).
Claim 3.3 easily follows from the following standard result by letting $\delta(a)=\operatorname{Rad}_{t}(a)$.
Theorem C. 3 (perturbation). Posit only one resource other than time (i.e., $d=2$ ). Consider a perturbation of the rescaled LP (C.1), where the reward vector $\boldsymbol{r}$ is replaced with $\tilde{\boldsymbol{r}}$, and the consumption vector $\boldsymbol{c}_{1}$ for the non-time resource is replaced with $\tilde{\boldsymbol{c}}_{1}$. Let $\tilde{\boldsymbol{X}}^{*}$ be its optimal solution. Assume $0 \leq \tilde{\boldsymbol{r}}-\boldsymbol{r} \leq \boldsymbol{\delta}$ and $0 \leq \boldsymbol{c}_{1}-\tilde{\boldsymbol{c}}_{1} \leq \boldsymbol{\delta}$, for some vector $\boldsymbol{\delta} \in[0,1]^{K}$. Then for each arm $a \neq a^{*}$,

$$
\delta(a)>G_{L A G}(a) \quad \text { if } \quad a \in \operatorname{supp}\left(\tilde{\boldsymbol{X}}^{*}\right)
$$

Proof. Let $\lambda_{1}^{*} \geq 0$ denote the dual variable corresponding to the single resource. Note that since $\mathrm{OPT}_{\mathrm{LP}} \leq 1$ and the dual vector $\boldsymbol{\lambda}^{*} \geq \mathbf{0}$ coordinate wise, we have $\lambda_{1}^{*} \leq 1$. From [18, Ch. 5.1] on local sensitivity when non-basic column of $A$ is changed, we have that the maximum allowable change to any single column $\delta(a) \leq \frac{\tilde{c}(a)}{\lambda_{1}^{*}}$ where $\tilde{c}(a)$ is the reduced-cost for the simplex algorithm, as defined in [18]. We will show that $\tilde{c}(a)=G_{\mathrm{LAG}}(a)$. Thus, if $\delta(a) \leq \frac{\tilde{c}(a)}{\lambda_{1}^{*}}=\frac{G_{\mathrm{LAG}}(a)}{\lambda_{1}^{*}}$ we have that the basis remains unchanged. Likewise from Bertsimas and Tsitsiklis [18, Ch. 5], the maximum allowed perturbation $\delta(a)$ on the reward $r(a)$ for the basis to remain unchanged is $\delta(a) \leq \tilde{c}(a)$. Combining these two we get the "if" part of the theorem.

It remains to prove that the reduced cost $\tilde{c}(a)=G_{\text {LAG }}(a)$. After converting the linear program to the standard form as required in [18], the reduced-cost $\tilde{c}(a)$ is given by the expression $\frac{T}{B\left(1-\eta_{\text {LP }}\right)} \sum_{j \in[d]} c_{j}(a) \tilde{\lambda}_{j}^{*}-r(a)$ where $\tilde{\lambda}^{*}$ is the optimal dual solution to LP (C.1). Note that $\boldsymbol{\lambda}^{*}:=\left(\frac{1}{1-\eta_{\mathrm{LP}}}\right) \tilde{\boldsymbol{\lambda}}^{*}$ is an optimal solution to the dual of the LP (2.2). Thus, plugging it into the definition of reduced cost and combining it with Claim C. 2 we have that

$$
\tilde{c}(a)=\frac{T}{B} \sum_{j \in[d]} \lambda_{j}^{*} c_{j}(a)-r(a)=G_{\mathrm{LAG}}(a)
$$

## D Various technicalities from Sections 3

## D. 1 Standard tools

We rely on some standard tools, which we state below for the sake of convenience.
Theorem D. 1 (Wald's identity). Let $X_{i}: i \in \mathbb{N}$ be i.i.d. real-valued random variables, adapted to filtration $\mathcal{F}_{i}: i \in \mathbb{N}$. Let $N$ be a stopping time relative to the same filtration. Then

$$
\mathbb{E}\left[X_{1}+X_{2}+\ldots+X_{N}\right]=\mathbb{E}\left[X_{i}\right] \cdot \mathbb{E}[N]
$$

Theorem D. 2 (Optimal Stopping Theorem). Let $X_{i}: i \in \mathbb{N}$ be a martingale sequence with $\mathbb{E}\left[X_{0}\right]=0$ adapted to filtration $\mathcal{F}_{i}: i \in \mathbb{N}$. Let $N$ be a stopping time relative to the same filtration. Then we have that $\mathbb{E}\left[X_{N}\right]=0$.
Theorem D. 3 ([37, 12]). Let $Z_{1}, Z_{2}, \ldots, Z_{T}$ be a martingale w.r.t. filtration $\left(\mathcal{F}_{t}\right)_{t \in[T]}$, such that $\left|Z_{t}\right| \leq c$ for all $t \in[T]$. Let $\mu:=\frac{1}{T} \sum_{t \in[T]} \mathbb{E}\left[Z_{t} \mid \mathcal{F}_{t-1}\right]$. Then,

$$
\operatorname{Pr}\left[\left|\sum_{t \in[T]} Z_{t}-\mu T\right|>\sqrt{2 \mu T c^{2} \ln \frac{T}{\delta}}\right] \leq \delta
$$

## D. 2 Proof of Eq. (3.6)

Let $\tau$ denote the stopping time of the algorithm that chooses arm $a^{*}$ in every time-step, given that the total budget is $B_{0}, T_{0}$ on the two resources. From definition we have $\operatorname{REW}\left(a^{*} \mid B_{0}, T_{0}\right)=$ $\sum_{t \in[\tau]} r_{t}\left(a^{*}\right)$. Using Wald's identity (Theorem D.1), we have that $\mathbb{E}\left[\operatorname{REW}\left(a^{*} \mid B_{0}, T_{0}\right)\right]=$ $\mathbb{E}[\tau] r\left(a^{*}\right)$.
Let $B_{0}, T_{0}$ denote the budget remaining for the two resources. By definition, we have that $\tau \geq T_{0}$ and $\sum_{t \in[\tau]} c_{t}\left(a^{*}\right) \geq B_{0}$. Using the Wald's identity (Theorem D.1) we have that $\mathbb{E}\left[\sum_{t \in[\tau]} c_{t}\left(a^{*}\right)\right]=$ $\mathbb{E}[\tau] c\left(a^{*}\right)$. Thus, we have $\mathbb{E}[\tau] \geq \min \left\{T_{0}, \frac{B_{0}}{c\left(a^{*}\right)}\right\} \geq \min \left\{T_{0}, B_{0}\right\}$. Therefore, we obtain the following.

$$
\begin{gather*}
\mathbb{E}\left[\operatorname{REW}\left(a^{*} \mid B_{0}, T_{0}\right)\right]=\mathbb{E}[\tau] r\left(a^{*}\right)>\left(\frac{\min \left\{T_{0}, B_{0}\right\}}{\max \left\{\frac{B}{T}, c\left(a^{*}\right)\right\}}\right) r\left(a^{*}\right), \quad \text { and }  \tag{D.1}\\
\mathbb{E}\left[\operatorname{REW}\left(a^{*} \mid B\right)\right]=\mathbb{E}\left[\tau_{B}\right] r\left(a^{*}\right) \leq\left(\frac{B}{\max \left\{\frac{B}{T}, c\left(a^{*}\right)\right\}}\right) r\left(a^{*}\right) \tag{D.2}
\end{gather*}
$$

Combining Equations (D.1) and (D.2), we get Eq. (3.6).

## D. 3 Proof of Eq. (3.9)

We now modify the above proof to get the tighter lower-bound in Eq. (3.9). Let $T_{0}, B_{0}$ denote the expected remaining time and budget (respectively) and let $\tau$ denote the (random) stopping time of the algorithm that chooses arm $a^{*}$ in every time-step given $T_{0}$ time-steps and $B_{0}$ budget. This implies that we have, $\mathbb{E}\left[\sum_{t \in[\tau]} c_{t}\left(a^{*}\right)\right] \geq B_{0}$ and $\mathbb{E}[\tau] \geq T_{0}$. From Theorem D.1, this implies that we have $\mathbb{E}[\tau] c\left(a^{*}\right) \geq B_{0}$ and $\mathbb{E}[\tau] \geq T_{0}$. This implies that $\mathbb{E}[\tau] \geq \min \left\{T_{0}, \frac{B_{0}}{c\left(a^{*}\right)}\right\}$.
Similar to Eq. (D.1) and Eq. (D.2) we obtain the following.

$$
\begin{gather*}
\mathbb{E}\left[\operatorname{REW}\left(a^{*} \mid B_{0}, T_{0}\right)\right]=\mathbb{E}[\tau] r\left(a^{*}\right)>\min \left\{T_{0}, \frac{B_{0}}{c\left(a^{*}\right)}\right\} r\left(a^{*}\right), \quad \text { and }  \tag{D.3}\\
\mathbb{E}\left[\operatorname{REW}\left(a^{*} \mid B_{0}=B, T_{0}=T\right)\right]=\mathrm{OPT}_{\mathrm{FD}} \leq\left(\frac{B}{\max \left\{\frac{B}{T}, c\left(a^{*}\right)\right\}}\right) r\left(a^{*}\right) . \tag{D.4}
\end{gather*}
$$

Combining Equations (D.3) and (D.4), we get Eq. (3.9).

## D. 4 Lower bound on Lagrange gap: Proof of Eq. (3.10)

We will use Eq. (3.4) and some standard properties of linear programming.
Assume $c\left(a^{*}\right)<\frac{B}{T}$. Using complementary slackness theorem on LP (2.2), this implies that $\lambda_{1}^{*}=0$. Moreover, note that the objective in the dual of LP (2.2) is $\lambda_{0}^{*}+\lambda_{1}^{*}=\lambda_{0}^{*}$. The optimal value of the primal LP (2.2) is $r\left(a^{*}\right)$ since, $X\left(a^{*}\right)=1$ is the optimal solution to the LP. This implies that $\lambda_{0}^{*}=r\left(a^{*}\right) \geq \frac{\mathrm{opT}_{\mathrm{FD}}}{T}$. Substituting this into Eq. (3.4) gives the first inequality in Eq. (3.10).
Now assume $c\left(a^{*}\right)>\frac{B}{T}$. Again, as above complementary slackness theorem on LP (2.2), this implies that $\lambda_{0}^{*}=0$. Thus, $G_{\mathrm{LAG}}(a)=\frac{T}{B} \cdot \lambda_{1}^{*} \cdot c(a)-r(a)$. Using the dual objective function $\lambda_{0}^{*}+\lambda_{1}^{*}=\lambda_{1}^{*}$ combined with strong duality, this implies that $\lambda_{1}^{*}=\frac{\mathrm{OPT}_{\mathrm{LP}}}{T} \geq \frac{\mathrm{OPT}_{\mathrm{FD}}}{T}$. Plugging this back into Eq. (3.4) gives the second inequality in Eq. (3.10).

## D. 5 Martingale arguments: Proof of Eq. (3.8)

For the proof of Eq. (3.8), we use the well-known theorem on optimal stopping time of martingales (Theorem D.2). Fix an arm $a \in[K]$. For any subset $S \subseteq[T]$ of rounds let $N_{S}(a), r_{S}(a)$ and $c_{S}(a)$ denote the number of times arm $a$ is chosen, the total realized rewards for arm $a$ and the total realized consumption of arm $a$, respectively. Let $\tau$ denote the (random) stopping time of a BwK algorithm with (random) budget $B$ and time $T$. Then we have the following claim.
Claim D.4. For a random stopping time $\tau$, for every arm $a \in[K]$ we have the following.

$$
\begin{align*}
& \mathbb{E}\left[r_{[\tau]}(a)\right]=r(a) \cdot \mathbb{E}\left[N_{[\tau]}(a)\right] .  \tag{D.5}\\
& \mathbb{E}\left[c_{[\tau]}(a)\right]=c(a) \cdot \mathbb{E}\left[N_{[\tau]}(a)\right] \tag{D.6}
\end{align*}
$$

Proof. We will prove the equality in Eq. (D.5); the one in Eq. (D.6) follows. Consider $r_{[\tau]}(a)$. By definition this is equal to $\sum_{t \in[\tau]} r_{t}(a) \cdot \mathbb{I}\left[a_{t}=a\right]$. Let $A_{t}:=\mathbb{I}\left[a_{t}=a\right]$ denote the random variable corresponding to the event that arm $a$ is chosen at time $t$. Define the random variable

$$
Y_{t}:=\sum_{t^{\prime} \leq t} A_{t^{\prime}} r_{t^{\prime}}(a)-\underset{t^{\prime}}{\mathbb{E}}\left[A_{t^{\prime}} r_{t^{\prime}}(a)\right]
$$

where $\mathbb{E}_{t}\left[\right.$.] denotes the conditional expectation given the random variables $A_{1}, A_{2}, \ldots, A_{t-1}$. It is easy to see that the sequences $\left\{X_{t}\right\}_{t \in[\tau]},\left\{Y_{t}\right\}_{t \in[\tau]}$ and $\left\{Z_{t}\right\}_{t \in[\tau]}$ forms a martingale sequence. Thus, we will apply the optimal stopping theorem (Theorem D.2) at time $\tau$, we have the following.

$$
\begin{equation*}
\mathbb{E}\left[Y_{\tau}\right]=\mathbb{E}\left[\sum_{t^{\prime} \leq \tau} A_{t^{\prime}} r_{t^{\prime}}(a)\right]-\mathbb{E}\left[\sum_{t^{\prime} \leq \tau} \underset{t^{\prime}}{\mathbb{E}}\left[A_{t^{\prime}} r_{t^{\prime}}(a)\right]\right]=0 \tag{D.7}
\end{equation*}
$$

Consider the term $\mathbb{E}\left[\sum_{t^{\prime} \leq \tau} \mathbb{E}_{t^{\prime}}\left[A_{t^{\prime}} r_{t^{\prime}}(a)\right]\right]$ in Eq. (D.7). This can be simplified to $\mathbb{E}\left[\sum_{t^{\prime} \leq \tau} r(a) \cdot \operatorname{Pr}\left[a_{t^{\prime}}=a\right]\right]$. Consider the following random variable

$$
Z_{t}:=\sum_{t^{\prime} \leq t} \operatorname{Pr}\left[a_{t^{\prime}}=a\right]-\underset{t^{\prime}}{\mathbb{E}}\left[\operatorname{Pr}\left[a_{t^{\prime}}=a\right]\right]
$$

Note that $\sum_{t^{\prime} \leq t} \mathbb{E}_{t^{\prime}}\left[\operatorname{Pr}\left[a_{t^{\prime}}=a\right]\right]=N_{[t]}(a)$. Thus, using Theorem D. 2 on the sequence $Z_{t}$ at the stopping time $\tau$, we obtain $\mathbb{E}\left[\sum_{t^{\prime} \leq \tau} \operatorname{Pr}\left[a_{t^{\prime}}=a\right]\right]=\mathbb{E}\left[N_{[\tau]}(a)\right]$.
Thus, the term $\mathbb{E}\left[\sum_{t^{\prime} \leq \tau} \mathbb{E}_{t^{\prime}}\left[A_{t^{\prime}} r_{t^{\prime}}(a)\right]\right]$ in Eq. (D.7) simplifies to $r(a) \cdot N_{[\tau]}(a)$ which gives the required equality in Eq. (D.5).

We will now use Claim D. 4 to prove Eq. (3.8). Recall that $\operatorname{REW}(a \mid B(a), T(a))$ denotes the total contribution to the reward by the BwK algorithm by playing arm $a$ with a (random) resource consumption of $B(a)$ and time steps of $T(a)$. Let $\tau$ be the (random) stopping time of this algorithm.

By definition we have that $N_{[\tau]}(a)=T(a)$. Thus, $\mathbb{E}\left[N_{[\tau]}(a)\right]=\mathbb{E}[T(a)$. From Eq. (D.6), we also have that $\mathbb{E}\left[N_{[\tau]}(a)\right]=\frac{\mathbb{E}\left[c_{[\tau \tau}(a)\right]}{c(a)}$. From the definition of $B(a)$ we have, $B(a)=c_{[\tau]}(a)$ and thus, $\mathbb{E}[B(a)]=\mathbb{E}\left[c_{[\tau]}(a)\right]$. Thus, this implies that $\mathbb{E}\left[N_{[\tau]}(a)\right]=\min \left\{T(a), \frac{\mathbb{E}[B(a)]}{c(a)}\right\}$.
Consider $\mathbb{E}[\operatorname{REW}(a)]=\mathbb{E}[\operatorname{REW}(a \mid B(a), T(a))]$.

$$
\begin{align*}
\mathbb{E}[\operatorname{REW}(a \mid B(a), T(a))] & =\mathbb{E}\left[r_{[\tau]}(a)\right] \\
& =r(a) \cdot \mathbb{E}\left[N_{[\tau]}(a)\right]  \tag{D.8}\\
& =r(a) \cdot \min \left\{T(a), \frac{\mathbb{E}[B(a)]}{c(a)}\right\}
\end{align*}
$$

$$
=r(a) \cdot \mathbb{E}\left[N_{[\tau]}(a)\right] \quad \text { (From Eq. (D.5)) }
$$

Now, consider $\operatorname{LP}(a \mid \mathbb{E}[B(a)], \mathbb{E}[T(a)])$. This value is equal to,

$$
\begin{aligned}
\mathbb{E}[\operatorname{REW}(a \mid \mathbb{E}[B(a)], \mathbb{E}[T(a)])] & =\frac{r(a)}{\max \{\mathbb{E}[B(a)] / \mathbb{E}[T(a)], c(a)\}} \cdot \frac{\mathbb{E}[B(a)]}{\mathbb{E}[T(a)]} \\
& =r(a) \cdot \min \left\{\mathbb{E}[T(a)], \frac{\mathbb{E}[B(a)]}{c(a)}\right\}
\end{aligned}
$$

Note that the last equality is same as the RHS in Eq. (D.8).

## E Proof of Theorem 3.10: generic $\sqrt{T}$ lower bound

Preliminaries. We rely on a well-known information-theoretic result for multi-armed bandits: essentially, no algorithm can reliably tell apart two bandit instances at time $T$ if they differ by at most $O(1 / \sqrt{T}) .{ }^{9}$ We formulate this result in a way that is most convenient for our applications.
Lemma E.1. Consider multi-armed bandits with Bernoulli rewards. Fix $\epsilon>0$ and two problem instances $\mathcal{I}, \mathcal{I}^{\prime}$ such that the mean reward of each arm differs by at most $\epsilon$ between $\mathcal{I}$ and $\mathcal{I}^{\prime}$. Suppose some bandit algorithm outputs distribution $\boldsymbol{Y}_{t}$ over arms at time $t \leq c / \epsilon^{2}$, for a sufficiently small absolute constant c. Let H be an arbitrary Lebesgue-measurable set of distributions over arms. Then either $\operatorname{Pr}\left[\boldsymbol{Y}_{t} \in H \mid \mathcal{J}_{t}=\mathcal{I}\right]>1 / 4$ or $\operatorname{Pr}\left[\boldsymbol{Y}_{t} \notin H \mid \mathcal{J}_{t}=\mathcal{I}^{\prime}\right]>1 / 4$ holds.

Applying Lemma E. 1 to bandits with knapsacks necessitates some subtlety. First, the rewards in the lemma will henceforth be called quasi-rewards, as they may actually correspond to consumption of a particular resource. Second, while a BwK algorithm receives multi-dimensional feedback in each round, the feedback other than the quasi-rewards will be the same (in distribution) for both problem instances, and hence can be considered a part of the algorithm. Third, distribution $\boldsymbol{Y}_{t}$ will be the conditional distribution over arms chosen by the BwK algorithm in round $t$ given the algorithm's observations so far; we will assume this without further mention. Fourth, we will need to specify the set $H$ of distributions (which will depend on a particular application).
Consider the rescaled LP (C.1) with $\eta_{\mathrm{LP}}:=6 * \mathrm{OPT}_{\mathrm{LP}} \sqrt{\frac{\log d T}{B}}$; we use this $\eta_{\mathrm{LP}}$ throughout this proof. Let $O T_{\mathrm{LP}}^{\mathrm{sc}}$ be the value of this LP. We prove the lower bound using $\mathrm{OPT}_{\mathrm{LP}}^{\mathrm{sc}}$ as a benchmark. This suffices by the following claim from prior work: ${ }^{10}$
Claim E. 2 (Immorlica et al. [35]). $O P T_{L P}^{s c} \leq O P T_{F D}$ for $\eta_{L P}:=6 \cdot O P T_{L P} \sqrt{\frac{\log d T}{B}}$.
Problem instances. Let $\boldsymbol{r}(a)$ and $\boldsymbol{c}(a) \in[0,1]^{d}$ be, resp., the mean reward and the mean resource consumption vector for each arm $a$ for instance $\mathcal{I}_{0}$. Let $\epsilon=c_{\mathrm{LB}} / \sqrt{T}$.
Problem instances $\mathcal{I}, \mathcal{I}^{\prime}$ are constructed as specified in the proof sketch; we repeat it here for the sake of convenience. For both instances, the rewards of each non-null arm $a \in\left\{A_{1}, A_{2}\right\}$ are deterministic and equal to $r(a)$. Resource consumption vector for arm $A_{1}$ is deterministic and equals $\boldsymbol{c}\left(A_{1}\right)$. Resource consumption vector of arm $A_{2}$ in each round $t$, denoted $\boldsymbol{c}_{(t)}\left(A_{2}\right)$, is a carefully

[^1]constructed random vector whose expectation is $c\left(A_{2}\right)$ for instance $\mathcal{I}$, and slightly less for instance $\mathcal{I}^{\prime}$. Specifically, $\boldsymbol{c}_{(t)}\left(A_{2}\right)=\boldsymbol{c}\left(A_{2}\right) \cdot W_{t} /\left(1-c_{\mathrm{LB}}\right)$, where $W_{t}$ is an independent Bernoulli random variable which correlates the consumption of all resources. We posit $\mathbb{E}\left[W_{t}\right]=1-c_{\mathrm{LB}}$ for instance $\mathcal{I}$, and $\mathbb{E}\left[W_{t}\right]=1-c_{\mathrm{LB}}-\epsilon$ for instance $\mathcal{I}^{\prime}$.

Main derivation. From the premise of the theorem (Eq. (3.15)), problem instance $\mathcal{I}$ admits an optimal solution $\boldsymbol{X}^{*}$ that is substantially supported on both non-null arms. Let $\boldsymbol{X}_{\mathcal{I}}^{*}, \boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}$ denote the optimal solutions to the scaled LP, instantiated for instances $\mathcal{I}, \mathcal{I}^{\prime}$ respectively.
The proof proceeds as follows. We first prove certain properties of distributions $\boldsymbol{X}_{\mathcal{I}}^{*}$ and $\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}$. We then use these properties and apply Lemma E. 1 with suitable quasi-rewards to complete the proof of the lower-bounds.
Since we modify the mean consumption of all resources for one arm in $\mathcal{I}^{\prime}$ this implies that $\boldsymbol{X}_{\mathcal{I}}^{*} \neq \boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}$. From assumption 3.8-(3.8) we have that $G_{\text {LAG }} \geq c_{\mathrm{LB}} / \sqrt{T}$. From the premise of the theorem, we have that the mean vector of consumptions for the resources $j \in[d]$ are all linearly independent. Thus, we can apply sensitivity theorem C. 3 to conclude that the support of the solution $\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}$ is same as $\boldsymbol{X}_{\mathcal{I}}^{*}$.
Moreover, from the linear independence of the consumption vectors and Eq. (3.15). combined with standard LP theory (see chapter 4 on duality in [18]) we have that there exists a resource $j^{*} \in[d]$ such that the optimal solution $\boldsymbol{X}_{\mathcal{I}}^{*}$ satisfies the resource constraint with equality.
In what follows, we denote the vector $\boldsymbol{c}$ as a shorthand for $\boldsymbol{c}_{j^{*}}$ (i.e., we drop the index $j^{*}$ ). Note that from the perturbation we have that $c\left(A_{1}\right)<c\left(A_{2}\right)$. Thus, for some $\delta>0$ we have $X_{\mathcal{I}^{\prime}}^{*}\left(A_{1}\right)=$ $X_{\mathcal{I}}^{*}\left(A_{1}\right)-\delta$ and $X_{\mathcal{I}^{\prime}}^{*}\left(A_{2}\right)=X_{\mathcal{I}}^{*}\left(A_{2}\right)+\delta$. Let $\|\boldsymbol{X}\|$ denote the $\ell_{1}$-norm of a given distribution $\boldsymbol{X}$. Thus, we have

$$
\begin{equation*}
\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}\right\|=2 \delta \tag{E.1}
\end{equation*}
$$

Given any distribution $\boldsymbol{Y}$ over the arms, let

$$
\begin{equation*}
V_{\mathrm{sc}}(\boldsymbol{Y}):=\left(1-\eta_{\mathrm{LP}}\right) \cdot B / T \cdot r(\boldsymbol{Y}) /\left(\max _{j \in[d]} c_{j}(\boldsymbol{Y})\right) . \tag{E.2}
\end{equation*}
$$

This is the value of $\boldsymbol{Y}$ in the rescaled LP (C.1), where $\boldsymbol{Y}$ itself is rescaled to make it LP-feasible (and as large as possible). Note that $V_{\mathrm{sc}}(\boldsymbol{Y})=\left(1-\eta_{\mathrm{LP}}\right) V(\boldsymbol{Y})$, where $V(\boldsymbol{Y})$ is the value of the original LP , as defined in (E.2). Also, $\mathrm{OPT}_{\mathrm{LP}}^{\mathrm{sc}}=\sup _{\boldsymbol{Y}} V_{\mathrm{sc}}(\boldsymbol{Y})$.
By a slight abuse of notation, let $V_{\mathrm{sc}}(\boldsymbol{Y}), V_{\mathrm{sc}}^{\prime}(\boldsymbol{Y})$ be the value of $V_{\mathrm{sc}}(\boldsymbol{Y})$ corresponding to instances $\mathcal{I}$ and $\mathcal{I}^{\prime}$ respectively.
We use the following two claims in the proof of our lower-bound. Claim E. 3 states that if a distribution is close to the optimal distribution for instance $\mathcal{I}$ then it is also far from the optimal distribution for $\mathcal{I}^{\prime}$. Claim E. 4 states that if a distribution is far from the optimal distribution, then playing from that distribution also incurs large instantaneous regret. Both claims have nothing to do with particular algorithms.

Claim E.3. Fix distribution $\boldsymbol{Y} \in \Delta^{3}$ and $\epsilon<1$. If $\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}\right\|<\epsilon \cdot c_{\mathrm{LB}}^{2}$ then $\left\|\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}-\boldsymbol{Y}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}$.
Claim E.4. Fix distribution $\boldsymbol{Y} \in \Delta^{3}$ and $\epsilon<1$. If $\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}$ then $V_{s c}\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-V_{s c}(\boldsymbol{Y}) \geq$ $\epsilon \cdot \frac{c_{\mathrm{CB}}^{3}}{2}$. Likewise, if $\left\|\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}-\boldsymbol{Y}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}$ then $V_{s c}^{\prime}\left(\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}\right)-V_{s c}^{\prime}(\boldsymbol{Y}) \geq \epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2}$.

We now invoke Lemma E. 1 with the quasi-rewards at each time-step determined by the consumption of the resource $j^{*}$.

Define the set,

$$
\begin{equation*}
\mathcal{H}:=\left\{\boldsymbol{Y}:\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}\right\} \tag{E.3}
\end{equation*}
$$

to complete the proof Theorem 3.10. Consider an arbitrary algorithm ALG. We consider two cases: $\mathcal{J}=\mathcal{I}$ and $\mathcal{J}=\mathcal{I}^{\prime}$, which denote the instance that satisfies the conclusion of this lemma for at least $\frac{T}{2}$ rounds for $T:=\frac{c_{\mathrm{LB}}}{\epsilon^{2}}$.

Let $\mathcal{J}=\mathcal{I}$. Let $\mathcal{T}$ denote the set of time-steps $t \in[T]$ such that $\mathcal{J}_{t}=\mathcal{I}$ and $\boldsymbol{Y}_{t} \in \mathcal{H}$. Then, the expected regret of ALG can be lower-bounded by,

$$
\begin{array}{rlr}
\mathbb{E}\left[\sum_{t \in \mathcal{T}} V_{\mathrm{sc}}\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-V_{\mathrm{sc}}\left(\boldsymbol{Y}_{t}\right)\right] & =\mathbb{E}\left[\sum_{t \in \mathcal{T}:\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}_{t}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}} V_{\mathrm{sc}}\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-V_{\mathrm{sc}}\left(\boldsymbol{Y}_{t}\right)\right] & \text { (by Eq. (E.3)) } \\
& \geq \mathbb{E}\left[\sum_{t \in \mathcal{T}} \epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2}\right] & \\
& \geq T / 4 \cdot \epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2} & \\
& \geq O\left(c_{\mathrm{LB}}^{4} \cdot \sqrt{T}\right) . & \text { (by Eq Lemma E. (E.4)) } \\
& \text { (Since } \left.\epsilon=\frac{c_{\mathrm{LB}}}{\sqrt{T}}\right)
\end{array}
$$

We use a similar argument when $\mathcal{J}=\mathcal{I}^{\prime}$. Let $\mathcal{T}^{\prime}$ denote the set of time-steps $t \in[T]$ such that $\mathcal{J}_{t}=\mathcal{I}^{\prime}$ and $\left\|\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}-\boldsymbol{Y}_{t}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}$. The expected regret of ALG can be lower-bounded by,

$$
\begin{array}{rlr}
\mathbb{E}\left[\sum_{t \in \mathcal{T}^{\prime}} V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}\right)-V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{Y}_{t}\right)\right] & =\mathbb{E}\left[\sum_{t \in \mathcal{T}^{\prime}:\left\|\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}-\boldsymbol{Y}_{t}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}} V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}\right)-V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{Y}_{t}\right)\right] & \\
& \geq \mathbb{E}\left[\sum_{t \in \mathcal{T}^{\prime}:\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}_{t}\right\|<\epsilon \cdot c_{\mathrm{LB}}^{2}} V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}\right)-V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{Y}_{t}\right)\right] & \\
& =\mathbb{E}\left[\sum_{t \in \mathcal{T}^{\prime}: \boldsymbol{Y}_{t} \notin \mathcal{H}} V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}\right)-V_{\mathrm{sc}}^{\prime}\left(\boldsymbol{Y}_{t}\right)\right] \text { Claim E.3) } \\
& \geq \mathbb{E}\left[\sum_{t \in[T]: \boldsymbol{Y}_{t} \notin \mathcal{H}} \epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2}\right] & \text { (by Eq. (E.3)) } \\
& \geq T / 4 \cdot \epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2} \\
& \geq O\left(c_{\mathrm{LB}}^{4} \cdot \sqrt{T}\right) . & \text { (by Eq. (E.4)) } \\
& & \text { (by Lemma E.1) } \\
& \text { (Since } \left.\epsilon=\frac{c_{\mathrm{LB}}}{\sqrt{T}}\right) .
\end{array}
$$

Proof of Claim E.3. Let $c\left(A_{1}\right), c\left(A_{2}\right)$ denote the expected consumption of arms $A_{1}$ and $A_{2}$ respectively in instance $\mathcal{I}$. Define $\zeta:=\frac{\epsilon c\left(A_{1}\right)}{1-c_{\mathrm{LB}}}$. By definition, this implies that the expected consumption of arm $A_{2}$ in instance $\mathcal{I}^{\prime}$ is $c\left(A_{2}\right)-\zeta$. Additionally, since the support contains two arms, we have that the following holds: $c\left(A_{1}\right) X_{\mathcal{I}}^{*}\left(A_{1}\right)+c\left(A_{2}\right) X_{\mathcal{L}}^{*}\left(A_{2}\right)=B / T *\left(1-\eta_{\mathrm{LP}}\right)$ and $c\left(A_{1}\right) X_{\mathcal{I}^{\prime}}^{*}\left(A_{1}\right)+c\left(A_{2}\right) X_{\mathcal{I}^{\prime}}^{*}\left(A_{2}\right)-\zeta X_{\mathcal{I}^{\prime}}^{*}\left(A_{2}\right)=B / T *\left(1-\eta_{\mathrm{LP}}\right)$. Thus, we have $c\left(A_{1}\right) X_{\mathcal{I}}^{*}\left(A_{1}\right)+c\left(A_{2}\right) X_{\mathcal{I}}^{*}\left(A_{2}\right)=c\left(A_{1}\right) X_{\mathcal{I}}^{*}\left(A_{1}\right)+c\left(A_{2}\right) X_{\mathcal{I}}^{*}\left(A_{2}\right)+\delta\left(C\left(A_{2}\right)-c\left(A_{1}\right)-\zeta\right)-\zeta X_{\mathcal{I}}^{*}\left(A_{2}\right)$.
Rearranging and using the assumptions in 3.8, we get that

$$
\begin{equation*}
\delta=\frac{\zeta X_{\mathcal{I}}^{*}\left(A_{2}\right)}{c\left(A_{2}\right)-c\left(A_{1}\right)-\zeta} \geq \frac{\epsilon c_{\mathrm{LB}}}{1-c_{\mathrm{LB}}} \cdot \frac{c_{\mathrm{LB}}}{1-2 c_{\mathrm{LB}}-\frac{\epsilon \cdot c_{\mathrm{LB}}}{1-c_{\mathrm{LB}}}} \geq \epsilon \cdot c_{\mathrm{LB}}^{2} \tag{E.4}
\end{equation*}
$$

Consider $\left\|\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}-\boldsymbol{Y}\right\|$. This can be rewritten as

$$
=\left\|\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}-\boldsymbol{Y}-\boldsymbol{X}_{\mathcal{I}}^{*}+\boldsymbol{X}_{\mathcal{I}}^{*}\right\|
$$

$$
\geq\left|\left\|\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}-\boldsymbol{X}_{\mathcal{I}}^{*}\right\|-\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}\right\|\right| \quad \text { (Triangle inequality) }
$$

$$
\geq 2 \delta-\epsilon \cdot c_{\mathrm{LB}}^{2} \quad \text { (Premise of the claim and Eq. (E.1)) }
$$

$$
\geq \epsilon \cdot c_{\mathrm{LB}}^{2} . \quad \quad \text { (From Eq. (E.4)) }
$$

Proof of Claim E.4. We will prove the statement $\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2} \Longrightarrow V_{\mathrm{sc}}\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-V_{\mathrm{sc}}(\boldsymbol{Y}) \geq$ $\epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2}$. The exact same argument holds by replacing $\boldsymbol{X}_{\mathcal{I}}^{*}$ with $\boldsymbol{X}_{\mathcal{I}^{\prime}}^{*}$ and $V_{\mathrm{sc}}($.$) with V_{\mathrm{sc}}^{\prime}($.$) .$

Consider $V_{\mathrm{sc}}\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-V_{\mathrm{sc}}(\boldsymbol{Y})$. By definition, this equals,

$$
\begin{equation*}
r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-\frac{r(\boldsymbol{Y})}{\max \left\{\frac{B^{\prime}}{T}, c(\boldsymbol{Y})\right\}} \cdot \frac{B^{\prime}}{T} \tag{E.5}
\end{equation*}
$$

where $B^{\prime}$ is the scaled budget.
We have two cases. In case 1 , let $\max \left\{\frac{B^{\prime}}{T}, c(\boldsymbol{Y})\right\}=\frac{B^{\prime}}{T}$. Thus, Eq. (E.5) simplifies to,

$$
\begin{aligned}
& =r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-r(\boldsymbol{Y}) \\
& =r\left(A_{1}\right)\left[X_{\mathcal{I}}^{*}\left(A_{1}\right)-Y\left(A_{1}\right)\right]+r\left(A_{2}\right)\left[X_{\mathcal{I}}^{*}\left(A_{2}\right)-Y\left(A_{2}\right)\right]
\end{aligned}
$$

Note that since $\max \left\{\frac{B^{\prime}}{T}, c(\boldsymbol{Y})\right\}=\frac{B^{\prime}}{T}$, this implies that $Y($ null $)=0$. Since $\boldsymbol{X}_{\mathcal{I}}^{*}$ is an optimal solution and $r\left(A_{2}\right)>r\left(A_{1}\right)$, this implies that we have $Y\left(A_{1}\right)=X_{\mathcal{I}}^{*}\left(A_{1}\right)+\zeta$ and $Y\left(A_{2}\right)=$ $X_{\mathcal{I}}^{*}\left(A_{2}\right)-\zeta$. Thus, we have,

$$
\begin{aligned}
r\left(A_{1}\right)\left[X_{\mathcal{I}}^{*}\left(A_{1}\right)-Y\left(A_{1}\right)\right]+r\left(A_{2}\right)\left[X_{\mathcal{I}}^{*}\left(A_{2}\right)-Y\left(A_{2}\right)\right] & \geq\left[r\left(A_{2}\right)-r\left(A_{1}\right)\right] \zeta \\
& \geq c_{\mathrm{LB}} \cdot\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}\right\| / 2 \\
& \geq \epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2} .
\end{aligned}
$$

Consider case 2 where $\max \left\{\frac{B^{\prime}}{T}, c(\boldsymbol{Y})\right\}=c(\boldsymbol{Y})$. Then, Eq. (E.5) simplifies to,

$$
\begin{aligned}
& =r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-\frac{B^{\prime}}{T} \cdot \frac{r(\boldsymbol{Y})}{c(\boldsymbol{Y})} \\
& \geq r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-\max _{\boldsymbol{Y} \in \Delta_{3}:\left\|\boldsymbol{X}_{\mathcal{I}}^{*}-\boldsymbol{Y}\right\| \geq \epsilon \cdot c_{\mathrm{LB}}^{2}} \frac{B\left(1-\eta_{\mathrm{LP}}\right)}{T} \cdot \frac{r(\boldsymbol{Y})}{c(\boldsymbol{Y})}
\end{aligned}
$$

The maximization happens when the distribution $\boldsymbol{Y}$ is such that $Y\left(A_{1}\right)=X_{\mathcal{I}}^{*}-\epsilon \cdot c_{\mathrm{LB}}^{2} / 2$ and $Y\left(A_{2}\right)=X_{\mathcal{I}}^{*}-\epsilon \cdot c_{\mathrm{LB}}^{2} / 2$. Plugging this into the expression we get the RHS is at least,

$$
\begin{aligned}
& \geq r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-\frac{B\left(1-\eta_{\mathrm{LP}}\right)}{T} \cdot \frac{r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)+\epsilon \cdot c_{\mathrm{LB}}^{2} / 2 \cdot\left(r\left(A_{2}\right)-r\left(A_{1}\right)\right)}{c\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)+\epsilon \cdot c_{\mathrm{LB}}^{2} / 2 \cdot\left(c\left(A_{2}\right)-c\left(A_{1}\right)\right)} \\
& \geq r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-c_{\mathrm{LB}}\left(1-\eta_{\mathrm{LP}}\right) \cdot \frac{r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)+\epsilon \cdot c_{\mathrm{LB}}^{2} / 2 \cdot\left(r\left(A_{2}\right)-r\left(A_{1}\right)\right)}{c\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)+\epsilon \cdot c_{\mathrm{LB}}^{2} / 2 \cdot\left(c\left(A_{2}\right)-c\left(A_{1}\right)\right)} \\
& \geq r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)-\left(1-\eta_{\mathrm{LP}}\right) \cdot \frac{r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right)+\epsilon \cdot c_{\mathrm{LB}}^{2} / 2 \cdot\left(r\left(A_{2}\right)-r\left(A_{1}\right)\right)}{1+\epsilon \cdot c_{\mathrm{LB}}^{2} / 2} \\
& \geq \frac{\eta_{\mathrm{LP}}}{2} \cdot r\left(\boldsymbol{X}_{\mathcal{I}}^{*}\right) \geq \epsilon \cdot \frac{c_{\mathrm{LB}}^{3}}{2} .
\end{aligned}
$$

The last two inequality follows from Assumption 3.8-(3.8), the value of $\eta_{\mathrm{LP}}$ and the fact that $\epsilon=\frac{c_{\mathrm{LB}}}{\sqrt{T}}$, respectively. Combining the two cases we get the claim.

## F Proof of Theorem 3.9(b): $\sqrt{T}$ lower bound for $d>2$

We first show that for any given instance $\mathcal{I}_{0}$, for a given $0<\delta_{1} \leq \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ we can obtain a $\delta_{1}$ perturbation of this instance, denoted by $\mathcal{I}_{0}^{\prime}$, that satisfies Eq. (3.15). Given instance $\mathcal{I}_{0}$ we construct the $\delta_{1}$-perturbation as follows. We construct instance $\mathcal{I}_{0}^{\prime}$ by decreasing the mean consumption on arm $A_{i}$ and resource $j$ by $\zeta_{1}^{j}$. We keep the mean rewards the same. Let $\boldsymbol{X}$ denote the optimal solution to instance $\mathcal{I}$. As a notation we denote the matrix $\boldsymbol{C} \in[0,1]^{d \times 3}$ as the matrix of mean consumption. Let $\boldsymbol{B}$ denote the sub-matrix of $\boldsymbol{C}$ such that, $\boldsymbol{X}$ satisfies the constraints in the scaled LP (C.1) with equality. Thus, we have $\boldsymbol{C} \cdot \boldsymbol{X}=\boldsymbol{b}$, where every co-ordinate of $b$ is $\frac{B\left(1-\eta_{\mathrm{L}}\right)}{T}$. Thus, the perturbation is equivalent to perturbing the vector $\boldsymbol{b}$, such that the $j^{t h}$ entry has an additive perturbation of $\zeta^{j}$. From Proposition 3.1 in [46], this linear program has a non-degenerate primal optimal solution, in the sense that it satisfies Eq. (3.15).
Next, we show that given an instance $\mathcal{I}_{0}^{\prime}$ we can obtain a $\delta_{2}$ perturbation of $\mathcal{I}_{0}^{\prime}$ for a given $0<$ $\delta_{2} \leq \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$, such that the consumption vectors are linearly independent. Define a random matrix
$\boldsymbol{D} \in\left[-\zeta_{2}, \zeta_{2}\right]^{d \times 3}$ such that every entry in $\boldsymbol{D}$ is generated uniformly at random from the set $\left[-\zeta_{2}, \zeta_{2}\right]$. We claim that the vectors $\boldsymbol{c}_{j}-\boldsymbol{d}_{j}$ are all linearly independent, where $\boldsymbol{d}_{j}$ is the $j^{\text {th }}$ row of $\boldsymbol{D}$ with probability at least 0.6 . In other words, decreasing each of the mean consumption by a uniformly random value chosen from the set $\left[-\zeta_{2}, \zeta_{2}\right]$ implies that there exists a realization of $\boldsymbol{D}$ such that the vectors $\boldsymbol{c}_{j}-\boldsymbol{d}_{j}$ are all linearly independent.

The proof of this claim proceeds as follows. As before define $C \in[0,1]^{d \times 3}$ to be the matrix of mean consumption. From definition of linear independence we need to show that the smallest singular value of the matrix $\boldsymbol{C}-\boldsymbol{D}$ is non-zero. Note that every entry in the matrix $\boldsymbol{C}-\boldsymbol{D}$ is chosen independently. Thus, using the bound on the probability of singularity in Theorem 2.2 of [21] we have that the probability that the smallest singular value is 0 is at most $\frac{1}{2 \sqrt{2}}$. Thus, with probability at least $1-\frac{1}{2 \sqrt{2}}>0.6$ we have that the matrix $\boldsymbol{C}-\boldsymbol{D}$ is singular.
Thus, for $\delta:=\delta_{1}+\delta_{2}$, we have that there exists a $\delta$-perturbed instance $\widetilde{\mathcal{I}}_{0}$, that satisfies all the assumptions in 3.8 and linear independence condition required in the premise of Theorem 3.10.

## G Simple regret: proof of Theorem 4.1

For convenience, let us restate the theorem:
Theorem. Assume $B \geq \Omega(T)$ and $\eta_{\mathrm{LP}} \leq \frac{1}{2}$. With probability at least $1-O\left(T^{-3}\right)$, for each $\epsilon>0$, there are at most $N_{\epsilon}=\mathcal{O}\left(\frac{K}{\epsilon^{2}} \log K T d\right)$ rounds $t$ such that $O P T_{D P} / T-r\left(\boldsymbol{X}_{t}\right) \geq \epsilon$.

The proof consists of two major steps: we argue about confidence sums, and we upper-bound simple regret in terms of the confidence radius.

## G. 1 Confidence sums

The following arguments depend only on the definition of the confidence radius, and work for any algorithm ALG. Suppose in each round $t$, this algorithm chooses a distribution $\boldsymbol{Y}_{t}$ over arms and samples arm $a_{t}$ independently $\boldsymbol{Y}_{t}$. We upper-bound the number of rounds $t$ with large $\operatorname{Rad}_{t}\left(\boldsymbol{Y}_{t}\right)$ :
Lemma G.1. Fix the threshold $\theta_{0}>0$, and let $S$ be the set of all rounds $t \in[T]$ such that $\operatorname{Rad}_{t}\left(\boldsymbol{Y}_{t}\right) \geq \theta_{0}$. Then $|S| \leq \mathcal{O}\left(\theta_{0}^{-2} \cdot K \log (K d T)\right)$ with probability at least $1-O\left(T^{-3}\right)$.

To prove the lemma, we study confidence sums: for a subset $S \subset[T]$ of rounds, define

$$
\begin{array}{lr}
W_{\text {act }}(S):=\sum_{t \in S} \operatorname{Rad}_{t}\left(a_{t}\right) & \text { (action-confidence sum of } A L G), \\
W_{\text {dis }}(S):=\sum_{t \in S} \operatorname{Rad}_{t}\left(\boldsymbol{Y}_{t}\right) & \text { (distribution-confidence sum of } A L G) .
\end{array}
$$

First, a standard argument (e.g., implicit in [10], see Section G.4) implies that

$$
\begin{equation*}
W_{\text {act }}(S) \leq \mathcal{O}\left(\sqrt{K|S| C_{\mathrm{rad}}}+K \cdot \ln |S| \cdot C_{\mathrm{rad}}\right) \quad \text { for any fixed subset } S \subset[T] \tag{G.1}
\end{equation*}
$$

Second, note that $W_{\text {dis }}(S)$ is close to $W_{\text {act }}(S)$ : for any fixed subset $S \subset[T]$,

$$
\begin{equation*}
\left|W_{\text {dis }}(S)-W_{\text {act }}(S)\right| \leq \mathcal{O}(\sqrt{|S| \log T}) \quad \text { with probability at least } 1-T^{-3} \tag{G.2}
\end{equation*}
$$

This is by Azuma-Hoeffding inequality, since $\left(\operatorname{Rad}_{t}\left(a_{t}\right)-\operatorname{Rad}_{t}\left(\boldsymbol{Y}_{t}\right): t \in S\right)$ is a martingale difference sequence. We extend this observation to random sets $S$. A random set $S \subset[T]$ is called time-consistent if the event $\{t \in S\}$ does not depend on the choice of arm $a_{t}$ or anything that happens afterwards, for each round $t$. (But it can depend on the choice of distribution $\boldsymbol{Y}_{t}$.)
Claim G.2. For any any time-consistent random set $S \subset[T]$,

$$
\begin{equation*}
\left|W_{\mathrm{dis}}(S)-W_{\mathrm{act}}(S)\right| \leq \mathcal{O}(\sqrt{|S| \log T}+\log T) \quad \text { with probability at least } 1-T^{-3} \tag{G.3}
\end{equation*}
$$

Proof. By definition of time-consistent set, for each round $t$,

$$
\mathbb{E}\left[\mathbf{1}_{\{t \in S\}} \cdot \operatorname{Rad}_{t}\left(a_{t}\right) \mid\left(\boldsymbol{Y}_{1}, a_{1}\right), \ldots,\left(\boldsymbol{Y}_{t-1}, a_{t-1}\right), \boldsymbol{Y}_{t}\right]=\mathbf{1}_{\{t \in S\}} \cdot \operatorname{Rad}_{t}\left(\boldsymbol{Y}_{t}\right)
$$

Thus, $\mathbf{1}_{\{t \in S\}} \operatorname{Rad}_{t}\left(a_{t}\right)-\operatorname{Rad}_{t}\left(\boldsymbol{Y}_{t}\right), t \in[T]$ is martingale difference sequence. Claim G. 2 follows from a concentration bound from prior work (Theorem D.3).

We complete the proof of Lemma G. 1 as follows. Fix $\delta>0$. Since $S$ is a time-consistent random subset of $[T]$, by Eq. (G.1) and Claim G.2, with probability at least $1-\delta$ it holds that

$$
\theta_{0} \cdot|S| \leq W_{\mathrm{dis}}(S) \leq \mathcal{O}\left(\sqrt{|S| K C_{\mathrm{rad}}}+K C_{\mathrm{rad}}+\sqrt{|S| \log T}+\log T\right)
$$

We obtain the Lemma by simplifying and solving this inequality for $|S|$.

## G. 2 Connecting LP-gap and the confidence radius

In what follows, let $B_{\mathbf{s c}}=B\left(1-\eta_{\mathrm{LP}}\right)$ be the budget in the rescaled LP.
Lemma G.3. Fix round $t \in[T]$, and assume the "clean event" in (2.7). Then

$$
G_{L P}\left(\boldsymbol{X}_{t}\right) \leq\left(2+T / B_{\mathrm{sc}}\right) \operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)
$$

Proof. Let $\alpha:=B_{\mathrm{sc}} / T$. For any distribution $\boldsymbol{X}$, let

$$
V_{+}(\boldsymbol{X}):=B_{\mathrm{sc}} / T \cdot r(\boldsymbol{X}) / \max _{j \in[d]} c_{j}^{-}(\boldsymbol{X})
$$

denote the value of $\boldsymbol{X}$ in the optimistic LP (2.6), after proper rescaling. Let $\boldsymbol{X}^{*}$ be an optimal solution to the (original) LP (2.2). Then

$$
\begin{equation*}
G_{\mathrm{LP}}\left(\boldsymbol{X}_{t}\right)=V\left(\boldsymbol{X}^{*}\right)-V\left(\boldsymbol{X}_{t}\right)-V_{+}\left(\boldsymbol{X}_{t}\right)+V_{+}\left(\boldsymbol{X}_{t}\right) \tag{G.4}
\end{equation*}
$$

Since $V_{+}\left(\boldsymbol{X}_{t}\right)$ is the optimal solution to the optimistic LP (2.6),

$$
V_{+}\left(\boldsymbol{X}_{t}\right) \geq V_{+}\left(\boldsymbol{X}^{*}\right)
$$

Moreover, since $\boldsymbol{X}^{*}$ is feasible to the optimistic LP (2.6) with the scaled budget $B_{\mathrm{sc}}$,

$$
V_{+}\left(\boldsymbol{X}^{*}\right) \geq V\left(\boldsymbol{X}^{*}\right)
$$

It follows that Eq. (G.4) an be upper-bounded as

$$
\begin{equation*}
G_{\mathrm{LP}}\left(\boldsymbol{X}_{t}\right) \leq V_{+}\left(\boldsymbol{X}_{t}\right)-V\left(\boldsymbol{X}_{t}\right) \tag{G.5}
\end{equation*}
$$

We will now upper-bound the right-hand side in the above. Denote

$$
\begin{aligned}
c_{\max }\left(\boldsymbol{X}_{t}\right) & :=\max _{j \in[d]} \sum_{a \in[K]} c_{j, t}(a) X_{t}(a) \\
c_{\max }^{-}\left(\boldsymbol{X}_{t}\right) & :=\max _{j \in[d]} \sum_{a \in[K]} c_{j, t}^{-}(a) X_{t}(a) .
\end{aligned}
$$

By definition of the value of a linear program, we can continue Eq. (G.5) as follows:

$$
\begin{align*}
G_{\mathrm{LP}}\left(\boldsymbol{X}_{t}\right) & \leq V_{+}\left(\boldsymbol{X}_{t}\right)-V\left(\boldsymbol{X}_{t}\right) \\
& \leq \alpha \cdot \frac{\hat{r}\left(\boldsymbol{X}_{t}\right)+\operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)}{c_{\max }^{-}\left(\boldsymbol{X}_{t}\right)}-\alpha \cdot \frac{r\left(\boldsymbol{X}_{t}\right)}{c_{\max }\left(\boldsymbol{X}_{t}\right)} \tag{G.6}
\end{align*}
$$

Under the clean event in Eq. (2.7), we continue Eq. (G.6) as follows:

$$
\begin{equation*}
\leq \alpha\left(\frac{2 \operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)+r\left(\boldsymbol{X}_{t}\right)}{c_{\max }^{-}\left(\boldsymbol{X}_{t}\right)}-\frac{r\left(\boldsymbol{X}_{t}\right)}{c_{\max }\left(\boldsymbol{X}_{t}\right)}\right) \tag{G.7}
\end{equation*}
$$

Since time is one of the resources, $c_{\max }^{-}\left(\boldsymbol{X}_{t}\right) \geq \frac{B_{\mathrm{sc}}}{T}$. Thus, we continue Eq. (G.7) as follows:

$$
\begin{align*}
& \leq 2 \operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)+\alpha r\left(\boldsymbol{X}_{t}\right)\left(\frac{1}{c_{\max }^{-}\left(\boldsymbol{X}_{t}\right)}-\frac{1}{c_{\max }\left(\boldsymbol{X}_{t}\right)}\right) \\
& =2 \operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)+\alpha r\left(\boldsymbol{X}_{t}\right)\left(\frac{\operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)}{c_{\max }^{-}\left(\boldsymbol{X}_{t}\right) \cdot c_{\max }\left(\boldsymbol{X}_{t}\right)}\right) \\
& \leq 2 \operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)+\frac{\operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)}{c_{\max }^{-}\left(\boldsymbol{X}_{t}\right)}  \tag{G.8}\\
& \leq\left(2+\frac{T}{B_{\mathrm{sc}}}\right) \operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right) \tag{G.9}
\end{align*}
$$

Eq. (G.8) uses the fact that $\alpha \frac{r\left(\boldsymbol{X}_{t}\right)}{c_{\max }\left(\boldsymbol{X}_{t}\right)} \leq \frac{B}{T} \frac{r\left(\boldsymbol{X}_{t}\right)}{c_{\max }\left(\boldsymbol{X}_{t}\right)}=V\left(\boldsymbol{X}_{t}\right) \leq 1$. Eq. (G.9) uses the fact that time is one of the resources and thus, $c_{\max }^{-}\left(\boldsymbol{X}_{t}\right) \geq \frac{B_{\mathrm{sc}}}{T}$.

## G. 3 Finishing the proof of Theorem 4.1

Claim G.4. Fix round $t$, and assume the "clean event" in (2.7). Then

$$
O P T_{D P} / T-r\left(\boldsymbol{X}_{t}\right) \leq G_{L P}\left(\boldsymbol{X}_{t}\right)+\eta_{\mathrm{LP}}
$$

Proof. By (2.7) and because $\boldsymbol{X}_{t}$ is the solution to the optimistic LP, we have

$$
\max _{j \in d} c_{j}\left(\boldsymbol{X}_{t}\right) \geq \max _{j \in d} c_{j}^{-}\left(\boldsymbol{X}_{t}\right)=B / T\left(1-\eta_{\mathrm{LP}}\right) .
$$

It follows that $r\left(\boldsymbol{X}_{t}\right) \geq V\left(\boldsymbol{X}_{t}\right)\left(1-\eta_{\mathrm{LP}}\right)$. Finally, we know that $\mathrm{OPT}_{\mathrm{LP}} \geq \mathrm{OPT}_{\mathrm{DP}} / T$.
Condition on (2.7), and the high-probability event in Lemma G.1. (Take the union bound in Lemma G. 1 over all thresholds $\theta_{0} \geq 1 / \sqrt{T}$, e.g., over an exponential scale.) Fix $\epsilon>0$. By Claim G. 4 and Lemma G.3, any round $t$ with simple regret at least $\epsilon$ satisfies

$$
\epsilon \leq \mathrm{OPT}_{\mathrm{DP}} / T-r\left(\boldsymbol{X}_{t}\right) \leq \eta_{\mathrm{LP}}+\left(2+T / B_{\mathrm{sc}}\right) \operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right)
$$

Therefore, $\operatorname{Rad}_{t}\left(\boldsymbol{X}_{t}\right) \geq \theta_{0}$, where $\theta_{0}=\frac{\epsilon-\eta_{\mathrm{LP}}}{\left(2+T / B_{\mathrm{sc}}\right)} \geq \Theta(\epsilon)$ when $\epsilon \geq 2 \eta_{\mathrm{LP}}$. Now, the theorem follows from Lemma G.1. Note, when $\epsilon<2 \eta_{\mathrm{LP}}$, then the total number of rounds in the theorem is larger than $T$ and hence not meaningful.

## G. 4 The standard confidence-sum bound: proof of Eq. (G.1)

Let us prove Eq. (G.1) for the sake of completeness. By definition of $\operatorname{Rad}_{t}\left(a_{t}\right)$ from Eq. (2.8),

$$
\operatorname{Rad}_{t}\left(a_{t}\right)=f(n):=\min \left(1, \sqrt{C_{\mathrm{rad}} / n}+C_{\mathrm{rad}} / n\right)
$$

where $N_{t}(a)$ is the number of times arm $a$ was chosen before round $t$. Therefore:

$$
\begin{aligned}
\sum_{t \in S} \operatorname{Rad}_{t}\left(a_{t}\right) & \leq \sum_{a \in[K]} \sum_{n=1}^{|S| / K} f(n) \\
& \leq \sum_{a \in[K]} \int_{x=1}^{|S| / K} f(x) \mathrm{d} x \leq 3\left(\sqrt{K|S| C_{\mathrm{rad}}}+K \cdot \ln |S| \cdot C_{\mathrm{rad}}\right)
\end{aligned}
$$

## H Reduction from BwK to bandits

We extend our results to any problem which can be cast as a special case of BwK and admits an upper bound on action-confidence sums, in the style of (G.1), for a suitably defined confidence radius.

To state the general result, let us define an abstract notion of "confidence radius". For each round $t$, a formal confidence radius is a mapping $\operatorname{Rad}_{t}(a)$ from algorithm's history and arm $a$ to $[0,1]$ such that with probability at least $1-O\left(T^{-4}\right)$ it holds that

$$
\left|r(a)-\hat{r}_{t}(a)\right| \leq \operatorname{Rad}_{t}(a) \quad \text { and } \quad\left|c_{j}(a)-\hat{c}_{j, t}(a)\right| \leq \operatorname{Rad}_{t}(a)
$$

for each resource $j$, where $\hat{r}_{t}(a)$ and $\hat{c}_{j, t}(a)$ denote average reward and resource consumption, as defined in Eq. (B.3). Such $\operatorname{Rad}_{t}(a)$ induces a version of UcbBwK with confidence bounds

$$
r_{t}^{+}(a)=\min \left(1, \hat{r}_{t}(a)+\operatorname{Rad}_{t}(a)\right) \quad \text { and } \quad c_{j, t}^{-}(a)=\max \left(0, \hat{c}_{j, t}(a)-\operatorname{Rad}_{t}(a)\right)
$$

We allow the algorithm to observe auxiliary feedback before and/or after each round, depending on a particular problem formulation, and this feedback may be used to compute the confidence radii.

We replace Eq. (G.1) with a generic bound on the action-confidence sum, for some $\beta$ that can depend on the parameters in the problem instance, but not on $S$ :

$$
\begin{equation*}
\sum_{t \in S} \operatorname{Rad}_{t}\left(a_{t}\right) \leq \sqrt{|S| \beta}, \quad \text { for any algorithm and any subset } S \subset[T] \tag{H.1}
\end{equation*}
$$

Theorem H.1. Consider an instance of BwK with time horizon T. Let $\operatorname{Rad}_{t}(\cdot)$ be a formal confidence radius which satisfies (H.1) for some $\beta$. Consider the induced algorithms UcbBwK and PrunedUcbBwK with rescaling parameter $\eta_{\mathrm{LP}}=\frac{2}{B} \sqrt{\beta T}$.
(i) Both algorithms obtain regret $O P T_{D P}-\mathbb{E}[R E W] \leq O(\sqrt{\beta T})\left(1+O P T_{D P} / B\right)$.
(ii) Theorem 3.2 holds with $\Psi=\beta G_{L A G}^{-2}$ and regret $\mathcal{O}\left(\beta G_{L A G}^{-1}\right)$ in part (ii).
(iii) Theorem 4.1 holds with $N_{\epsilon}=\mathcal{O}\left(\beta \epsilon^{-2}\right)$.

Proof Sketch For part (i), the analysis in [3] explicitly relies on (G.1). For part (ii), we modify the proof of Theorem 3.2 so as to use (G.1) instead of Claim 3.4. For part (iii), our proof of Theorem 4.1 uses (G.1) explicitly. In all three parts, we replace (G.1) with (H.1), and trace how the latter propagates through the respective proof.

We apply this general result to three specific scenarios: linear contextual bandits with knapsacks (LinCBwK) [5], combinatorial semi-bandits with knapsacks (SemiBwK) [49], and multinomial-logit bandits with knapsacks (MnlBwK) [26]. In all three applications, the confidence-sum bound (H.1) is implicit in prior work on the respective problem without resources. The guarantees in part (i) match those in prior work referenced above, up to logarithmic factors, and are optimal when $B=\Omega(T)$; in fact, we obtain an improvement for MnlBwK. Parts (ii) and (iii) - the results for logarithmic regret and simple regret - did not appear in prior work.

## H. 1 Linear Contextual Bandits with Knapsacks (LinCBwK)

In Contextual Bandits with Knapsacks (CBwK), we have $K$ actions, $d$ resources, budget $B$ and time horizon $T$, like in BwK, and moreover we have a set $\mathcal{X}$ of possible contexts. At each round $t \in[T]$, the algorithm first obtains a context $\boldsymbol{x}_{t} \in X$. The algorithm then chooses an action $a_{t} \in[K]$ and obtains an outcome $\boldsymbol{o}_{t}\left(a_{t}\right) \in[0,1]^{d+1}$ like in $\operatorname{BwK}$. The tuple ( $\left.\boldsymbol{x}_{t} ; \boldsymbol{o}_{t}(a): a \in[K]\right)$ is drawn independently from some fixed but unknown distribution. The algorithm continues until some resource, including time, is exhausted. One compares against a given a set $\Pi$ of policies: mappings from contexts to actions. We can formally interpret CBwK as an instance of BwK in which actions correspond to policies in $\Pi$. This interpretation defines the benchmarks $O P_{D P}$ and $O \mathrm{PT}_{\mathrm{FD}}$ that we compete with.
LinCBwK is a special case of $\operatorname{CBwK}$ in which the context space is $\mathcal{X}=[0,1]^{K \times m}$, for some parameter $m \in \mathbb{N}$, so that each context $\boldsymbol{x}_{t}$ is in fact a tuple $\boldsymbol{x}_{t}=\left(\boldsymbol{x}_{t}(a) \in[0,1]^{m}: a \in[K]\right)$. We have a linearity assumption: for some unknown matrix $\boldsymbol{W}_{*} \in[0,1]^{m \times(d+1)}$ and each arm $a \in[K]$,

$$
\mathbb{E}\left[\boldsymbol{o}_{t}(a) \mid \boldsymbol{x}_{t}(a)\right]=\boldsymbol{W}_{*}^{\mathrm{T}} \cdot \boldsymbol{x}_{t}(a)
$$

The policy set $\Pi$ consists of all possible policies.
Linear contextual bandits, studied in prior work [e.g., 9, 29, 43, 27, 2], is the special case without resources. Much of the complexity of linear contextual bandits (resp., LinCBwK) is captured by the special case of of linear bandits (resp., linear BwK) where the context is the same in each round.

The general theme in the work on linear bandits (contextual or not) to replace the dependence on the number of arms $K$ in the regret bound with the dependence on the dimension $m$ and, if applicable, avoid the dependence on $|\Pi|$. This is what we accomplish, too.
Corollary H.2. For LinCBwK, Theorem H.l holds with $\beta=\mathcal{O}\left(m^{2} d^{2} \log (m T d)\right)$.
Proof. Combining Lemma 13 of [9] and Theorem 2 of [1], it follows that the confidence-sum bound Eq. (H.1) holds with $\beta=\mathcal{O}\left(m^{2} d^{2} \log m T d\right)$.

## H. 2 Combinatorial Semi-bandits with Knapsacks (SemiBwK)

SemiBwK is a version of BwK, where actions correspond to subsets of some fixed ground set $[N]$ (whose elements are called atoms). There is a fixed family $\mathcal{F} \subset 2^{[N]}$ of feasible actions. In each round $t$, the algorithm chooses a subset $A_{t} \in \mathcal{F}$ and observes the outcome $\boldsymbol{o}_{t}(a) \in[0,1 / n]^{d}$ for each atom $a \in A_{t}$, where $n=\max _{A \in \mathcal{F}}|A|$. The outcome for a given subset $A \in \mathcal{F}$ is defined as the sum

$$
\begin{equation*}
\boldsymbol{o}_{t}(A)=\sum_{a \in A} \boldsymbol{o}_{t}(a) \in[0,1]^{d+1} \tag{H.2}
\end{equation*}
$$

The outcome matrix $\left(o_{t}(a): a \in[N]\right)$ is drawn independently from some fixed but unknown distribution. The algorithm continues until some resource, including time, is exhausted.
Combinatorial semi-bandits, the problem studied in prior work [e.g., 25, 40, 39], is the special case without resources. Note that the number of feasible actions can be exponential in $N$. The general
theme in this line of work is to replace the dependence on $|\mathcal{F}|$ in the regret bound with the dependence on $N$, or, even better, on $n$. We extend this to SemiBwK.

Corollary H.3. For SemiBwK, Theorem H. 1 holds with $\beta=\mathcal{O}(n \log (N d T))$.

Proof. Using Lemma 4 in [60] we immediately obtain the confidence-sum bound Eq. (H.1) with $\beta=n \log K d T$.

## H. 3 Multinomial-logit Bandits with Knapsacks (Mn1BwK)

In the MnlBwK problem, the setup starts like in SemiBwK. There is a ground set of $N$ atoms, and a fixed family $\mathcal{F} \subset 2^{[N]}$ of feasible actions. In each round, each atom $a$ has an outcome $\boldsymbol{o}_{t}(a) \in[0,1]^{d+1}$, and the outcome matrix $\left(\boldsymbol{o}_{t}(a): a \in[N]\right)$ is drawn independently from some fixed but unknown distribution. The aggregate outcome is formed in a different way: when a given subset $A_{t} \in \mathcal{F}$ is chosen by the algorithm in a given round $t$, at most one atom $a_{t} \in A_{t}$ is chosen stochastically by "nature", and the aggregate outcome is then $\boldsymbol{o}_{t}\left(A_{t}\right):=\boldsymbol{o}_{t}(a)$; otherwise, the algorithm skips this round. A common interpretation is that the atoms correspond to products, the chosen action $A_{t} \in \mathcal{F}$ is the bundle of products offered to the customer, and at most one product from this bundle is actually purchased. As usual, the algorithm continues until some resource (incl. time) is exhausted.
The selection probabilities are defined via the multinomial-logit model. For each atom $a$ there is a hidden number $v_{a} \in[0,1]$, interpreted as the customers' valuation of the respective product, and the

$$
\operatorname{Pr}\left[\text { atom } a \text { is chosen } \mid A_{t}\right]= \begin{cases}\frac{v_{a}}{1+\sum_{a^{\prime} \in A_{t}} v_{a^{\prime}}} & \text { if } a \in A_{t} \\ 0 & \text { otherwise }\end{cases}
$$

The set $\mathcal{F}$ of possible bundles is

$$
\mathcal{F}=\{A \subset[N]: \boldsymbol{M} \cdot x(A) \leq \boldsymbol{b}\}
$$

for some (known) totally unimodular matrix $\boldsymbol{M} \in \mathbb{R}^{N \times N}$ and a vector $\boldsymbol{b} \in \mathbb{R}^{N}$, where $x(A) \in$ $\{0,1\}^{N}$ represents set $A$ as a binary vector over atoms.

Multinomial-logit bandits, the problem studied in prior work [e.g., 7, 48, 51, 24], is the special case without resources. We derive the following corollary from the analysis of MNL-bandits in Agrawal et al. [7], which analyzes the confidence sum for the $v_{a}$ 's.
Corollary H.4. Consider MnlBwK and denote $V:=\sum_{a \in[N]} v_{a}$. Theorem H. 1 holds with

$$
\beta=\mathcal{O}\left(\left(\frac{\ln T}{\ln (1+1 / V)}\right)^{2}(N \sqrt{\ln (N T)}+\ln (N T))\right)=\widetilde{O}\left(N^{3}\right)
$$

Proof. The proof is implicit in the analysis in Agrawal et al. [7]. As in their paper, let $n_{\ell}$ denote the number of time-steps in phase $\ell$. Let $V_{\ell}=\sum_{a \in S_{\ell}} v_{a}$. Recall that $n_{\ell}$ is a geometric random variable with mean $\frac{1}{1+V_{\ell}}$. Using Chernoff-Hoeffding bounds we obtain that with probability at least $1-\frac{1}{T^{2}}$, $n_{\ell} \leq \frac{\ln T}{\ln \left(1+1 / V_{\ell}\right)}$.
Consider a random subset $S$. Summing the LHS and RHS in Lemma 4.3, we get that $\sum_{t \in S} \operatorname{Rad}_{t}\left(a_{t}\right) \leq \sum_{a \in[N]} \sum_{\ell: t \in \mathcal{T}_{a}(\ell)} \tilde{R}_{a}\left(S_{\ell}\right)$. Using Lemma 4.3 in [7] we have, $\sum_{a \in[N]} \sum_{\ell: t \in \mathcal{T}_{a}(\ell)} \tilde{R}_{a}\left(S_{\ell}\right) \leq \sum_{a \in[N]} \sum_{\ell: t \in \mathcal{T}_{a}(\ell)} n_{\ell} \sqrt{\frac{v_{a} \ln \sqrt{N} T}{T_{a}(\ell)}}+\frac{\ln \sqrt{N} T}{T_{a}(\ell)}$. Note that $v_{a} \leq 1$. Using the upper bound on $n_{\ell}$ derived above combined with the argument used to obtain (A.19) in [7] we get the desired value of $\beta$.

The worst-case regret bound from Corollary H. 4 improves over prior work [26]. In particular, consider the worst-case dependence on $N$, the number of atoms. Our regret bound scales as $N^{3 / 2}$, whereas the regret bound in [26] scales as $N^{7 / 2}$ (while both scale as $\sqrt{T}$ ).

## H. 4 Computational issues

We do not provide a generic computationally efficient implementation for UcbBwK in our reduction. The algorithm constructs and solves a linear program in each round, with one variable per arm in the reduction. So, even if the regret is fairly small, the number of LP variables may be very large: indeed, it may be exponential in the number of atoms in SemiBwK and MnlBwK, arbitrarily large compared to the other parameters in linear BwK, or even infinite as in LinCBwK. The corresponding LPs have a succinct representation in all these applications, but we do not provide a generic implementation. However, such (or very similar) linear programs may be computationally tractable via applicationspecific implementations, and indeed this is the case in LinCBwK [5] and SemiBwK [49]. In the prior work on MnlBwK [26], the $\sqrt{T}$-regret algorithm is not computationally efficient, same as ours; there is, however, a computationally efficient algorithm with regret $T^{2 / 3}$.


[^0]:    ${ }^{8}$ This framing with raw material(s) — BwK formulations of revenue management problems in which products being sold are separate from raw material(s) being consumed - traces back to Besbes and Zeevi [20].

[^1]:    ${ }^{9}$ This strategy for proving lower bounds in multi-armed bandits goes back to Auer et al. [11]. Lemma E. 1 is implicit in Auer et al. [11], see Slivkins [54, Lemma 2.9] for exposition.
    ${ }^{10}$ Claim E. 2 is a special case of Lemma 8.6 in Immorlica et al. [35] for $\tau^{*}=T$ and the reward/consumption for each arm, each resource and each time-step replaced with the mean reward/consumption.

