

Better Algorithms for Stochastic Bandits with Adversarial Corruptions

Anupam Gupta

Department of Computer Science, Carnegie Mellon University, Pittsburgh PA, USA

ANUPAMG@CS.CMU.EDU

Tomer Koren

Kunal Talwar

Google, Mountain View CA, USA

TKOREN@GOOGLE.COM

KUNAL@GOOGLE.COM

Editors: Alina Beygelzimer and Daniel Hsu

Abstract

We study the stochastic multi-armed bandits problem in the presence of adversarial corruption. We present a new algorithm for this problem whose regret is nearly optimal, substantially improving upon previous work. Our algorithm is agnostic to the level of adversarial contamination and can tolerate a significant amount of corruption with virtually no degradation in performance.

1. Introduction

We consider the stochastic multi-armed bandit problem with adversarial corruptions. The classical stochastic multi-armed bandits problem ([Lai and Robbins, 1985](#)) has long been studied, with many algorithms including the Upper Confidence Bound (UCB) algorithm ([Auer et al., 2002a](#)) and the Active Arm Elimination (AAE) algorithm ([Even-Dar et al., 2006](#)). (See, e.g., [Bubeck et al., 2012](#) for more details and many references.) Since the assumption that the rewards of the arms are stochastic might appear quite rigid, the *adversarial* multi-armed bandit model has also been studied, with algorithms like the EXP3 algorithm of [Auer et al. \(2002b\)](#).

The robustness of adversarial multi-armed bandit algorithms, however, comes at a heavy price: their guarantees are often significantly weaker than those of stochastic algorithms, the regret of which scales only logarithmically with the number of decision rounds. A natural question is therefore: can we obtain results with guarantees that degrade smoothly as one goes from the stochastic setting towards the adversarial one? More concretely, we can retain the favorable guarantees of the stochastic setting while tolerating a small amount of adversarial corruption? These questions were the focus of the recent work by [Lykouris et al. \(2018\)](#), who studied the problem of stochastic multi-armed bandit with adversarial corruptions. They motivated it by a variety of practical scenarios: these range from click fraud in pay-per-click advertisements, to the presence of spam and malicious reviews in recommendation systems.

The model of [Lykouris et al. \(2018\)](#) is a variant of the classic stochastic multi-armed bandit problem, where each pull of an arm generates a stochastic reward that may be contaminated by an adversary before it is revealed to the player. The difficulty of the problem then scales with the total amount of corruption introduced by the adversary. [Lykouris et al. \(2018\)](#) point out that standard stochastic algorithms like UCB and AAE can suffer severely from corruptions because the adversary, using a small amount of corruption, can cause the actual best arm to be eliminated by the algorithm, and hence make it fail miserably. To address this, they show how to add robustness to these algorithms at the cost of moderate degradation in performance.

1.1. The model

Before proceeding, let us describe the setup in some more detail. A player is faced with K arms, indexed by $[K] := \{1, 2, \dots, K\}$. Each arm $i \in [K]$ is associated with a reward distribution, unknown to the player, with mean μ_i . The best arm $i^* = \arg \max_i \mu_i$ is the one that maximizes the expected reward; let $\mu^* = \mu_{i^*}$ denote the maximal expected reward, and for any suboptimal arm $i \neq i^*$ consider its gap $\Delta_i = \mu^* - \mu_i \in [0, 1]$. The player interacts with the arms and has to repeatedly decide on an arm to pull over the course of T decision rounds. This interaction is contaminated by an adversary, that may corrupt the feedback observed by the player on at most C of these T rounds. (The eventual definition of C will be a bit more general.) After each arm pull, the player observes the possibly corrupted reward corresponding to that arm, and nothing else; this is traditionally called “bandit feedback.” The player’s goal is to minimize her regret, which is the difference between the total reward of the best arm and her own cumulative reward.

In this setting, the algorithm of [Lykouris et al. \(2018\)](#) achieves a regret bound of the form $\tilde{O}(KC \sum_{i \neq i^*} 1/\Delta_i)$.¹ That is, their bound is $O(KC)$ times worse than the standard $\tilde{O}(\sum_{i \neq i^*} 1/\Delta_i)$ bound achievable in the uncorrupted case by algorithms for the stochastic case such as UCB, yet is still poly-logarithmic in the number of rounds T . Importantly, their algorithm is *agnostic* to the level of corruption inflicted by the adversary and need not know the value of C ahead of time. It is based on a “multi-layered” scheme that automatically adapts to the amount of adversarial contamination: the first layer is the fastest but also the most susceptible to corruptions, whereas the last is the most robust but is also the slowest to learn.

1.2. Our results

We show how to get significantly better regret bounds for stochastic multi-armed bandits with adversarial corruptions. We give a new algorithm for the problem that attains a regret bound which removes the multiplicative dependence on C appearing in the existing bounds and trades it for an additive dependence. Our algorithm is entirely agnostic to the corruption level C , and is also extremely simple. Formally, our main theorem is the following:

Theorem (Theorem 1, informal). *There exists an algorithm for stochastic MAB with adversarial corruptions (called BARBAR; see [Algorithm 1](#) below), the regret of which is with high probability at most*

$$O(KC) + \tilde{O}\left(\sum_{i \neq i^*} \frac{1}{\Delta_i}\right).$$

To put the parameters in context, consider an instance with K arms and gaps $\Delta_i = \sqrt{K/T}$; in the uncorrupted case, this is a prototypical hard instance where the minimax regret is $\tilde{\Theta}(\sqrt{KT})$ and is matched by both stochastic algorithms, like UCB, and adversarial algorithms, like EXP3. (In fact, this instance was used to show the $\Omega(\sqrt{KT})$ lower bound for adversarial bandits by [Auer et al., 2002b](#).) In the corrupted case, however, algorithms for the stochastic case might fail and suffer linear $\Omega(T)$ regret. We are robust to corruptions up to the level $C = O(\sqrt{T/K})$, obtaining essentially the same regret as in the uncorrupted case. In contrast, the previous bounds from [Lykouris et al. \(2018\)](#) can only guarantee $O(K^{1.5}C\sqrt{T})$ regret, and thus suffer from a degradation in performance due to the corruptions already at level $C = O(1)$ while becoming vacuous when $C = \Omega(\sqrt{T/K})$.

In the uncorrupted case, where $C = 0$, our result recovers the classical bounds (achieved, e.g., by the UCB algorithm) up to a logarithmic factor. We also show that under several different conditions,

1. Here and throughout, the \tilde{O} notation suppresses all dependence on logarithmic terms.

it is possible to improve our regret bound and remove the factor of K from its first term; we leave as an open question whether this is possible to achieve in the general, agnostic case without any further assumptions.

1.3. Related work

Bridging between the stochastic and adversarial settings in multi-armed bandits, and in online learning more generally, has been a topic of significant interest in recent years. Most research in this direction has focused on obtaining “the best of both worlds” guarantees (Bubeck and Slivkins, 2012; Seldin and Slivkins, 2014; Auer and Chiang, 2016; Seldin and Lugosi, 2017; Zimmert and Seldin, 2018; Zimmert et al., 2019). The goal there is to achieve the better of the bounds at the two extremes: the worst case $O(\sqrt{T})$ -type bound on any problem instance, and the better $O(\log T)$ -type bound whenever the instance is actually stochastic. In particular, algorithms of this kind cannot adapt to instances that are close to being stochastic, which are treated as fully adversarial. The model we consider here aims to gradually interpolate between the two extremes, as one goes from the stochastic setting and injects a moderate amount of adversarial corruptions.

Adversarial contaminations similar to those considered here have been studied before in the context of bandit problems. Seldin and Slivkins (2014) and Zimmert and Seldin (2018) consider a “moderately contaminated” regime in which the adversarial corruptions do not reduce the gap $\min_{i \neq i^*} \Delta_i$ by more than a constant factor at any point in time. The latter paper implies strong regret bounds, comparable to the bounds given in this paper, but relies on a stringent condition that constrains the leading arm (i.e., the one with the highest empirical cumulative reward) at any point in time to being the true best arm. (We remark that our general bound can also be improved under various conditions, see Section 4.) Kapoor et al. (2018) study a closely related model, where each step is contaminated independently with probability η . They also study a contextual bandit model where the corruption is adversarial, subject to the constraint that for any prefix, at most an η fraction of the steps are corrupted. Altschuler et al. (2018) consider contaminated stochastic bandits in the closely related best-arm identification setting, where rounds of a stochastic bandits instance are contaminated uniformly at random by an “outlier” distribution, and the goal is to identify the best arm in spite of the corruptions. This model of contamination is again weaker than the adaptive adversarial model we consider here, and their goal is also somewhat different than ours.

A somewhat different flavor of robustness in multi-armed bandits has been explored by Bubeck et al. (2013). Their “Robust UCB” algorithm can tolerate heavy-tailed distributions of rewards; namely, it does not require boundedness or sub-Gaussianity of the rewards, and instead only needs them to have bounded variance (or any higher-order moment). These ideas were later adapted to other bandit settings (e.g., Yu et al., 2018; Shao et al., 2018).

Finally, we acknowledge the long line of research in statistics (Huber, 1964) and statistical (batch) learning (Valiant, 1985; Kearns and Li, 1988) on making algorithms robust to various kinds of adversarial perturbations. There have been several recent advances in computationally efficient algorithms for problems of this kind; see, e.g., Diakonikolas et al. (2016); Lai et al. (2016); Charikar et al. (2017); Diakonikolas et al. (2018); Klivans et al. (2018) and the references therein.

2. Setup: Stochastic MAB with Adversarial Corruptions

We consider a multi-armed bandit problem with K arms, indexed by $[K] = \{1, 2, \dots, K\}$. Each arm $i \in [K]$ is associated with an unknown reward distribution with mean μ_i ; for simplicity we assume that the reward distributions are all supported in $[0, 1]$, in which case $\mu_i \in [0, 1]$. We let

$i^* = \arg \max_i \mu_i$ denote the best arm, and let $\mu^* = \mu_{i^*}$ be the maximal expected reward. Finally, for any arm $i \neq i^*$ we define the gap $\Delta_i = \mu^* - \mu_i \in [0, 1]$; to simplify notation we assume that the optimal arm is unique, in which case $\Delta_i > 0$ for all $i \neq i^*$.

A player has to repeatedly choose an arm to pull and observe its reward as feedback, that may be contaminated by an adversary; without loss of generality, we assume that the adversary is deterministic. Formally, the interaction proceeds as follows. On each decision round $t = 1, \dots, T$:

- (i) A stochastic reward $R_i^t \in [0, 1]$ is drawn for each arm i from its reward distribution;
- (ii) The adversary observes the vector R^t and generates corrupted rewards $\tilde{R}^t \in [0, 1]^K$;
- (iii) The player picks $i_t \in [K]$, possibly at random, and observes the corrupted reward $\tilde{R}_{i_t}^t$.

The player is evaluated using the standard metric in stochastic MAB models of *pseudo-regret* (henceforth called simply *regret*; see remark below about other possible notions of regret), given by

$$\mathcal{R}(T) = \sum_{t=1}^T (\mu^* - \mu_{i_t}).$$

We will be interested in bounding the player's regret in terms of the level of corruption introduced by the adversary. To this end, we measure the level of corruption C as:

$$C = \sum_{t=1}^T \|\tilde{R}^t - R^t\|_\infty.$$

Our main focus will be on the case where C is unknown, and the player has to be *agnostic* to the level of corruption.

Remarks. Notice that we allow the adversary to be adaptive, in the sense that the corruptions on round t may be determined as a function of the player's past choices as well as of the stochastic rewards in the current and previous rounds. (However, the corruption on round t is independent of the choice of the player on the same round.) As a consequence, the corruption level C is a random variable that depends on the stochastic rewards as well as on the player's internal randomization; the bounds we provide in this paper will depend on the realized value of this random variable.

Also note that our definition of regret above is precisely the familiar pseudo-regret metric ubiquitous in stochastic MAB models. Our results directly extend to several other notions of regret; we discuss this in detail in [Section 5](#).

3. The BARBAR Algorithm

At a high level, our algorithm is similar in spirit to the Active Arm Elimination (AAE) algorithm for stochastic bandits, but contains some crucial modifications that make it robust to corruptions.

The algorithm proceeds in epochs which increase exponentially in length: the m^{th} epoch has length roughly 2^{2m} . At the beginning of each epoch m , the algorithm computes an empirical estimate Δ_i^{m-1} of the gap of each arm i based on its pulls during the previous epoch; namely, we take the difference between the reward-per-pull for arm i and a lower confidence bound for the highest reward-per-pull of any arm in epoch $m-1$. The algorithm uses these estimates to bias towards pulling the seemingly-better arms: each arm is pulled roughly $(\Delta_i^{m-1})^{-2}$ times, but never more than 2^{2m} times, during epoch m . As a consequence, the algorithm never eliminates an arm permanently; it

gives the seemingly-not-so-good arms some recourse by pulling them a small number of times in each epoch. This, along with the fact that each epoch only uses information from the immediately preceding epoch, ensures that any corruption has a bounded impact on the algorithm's regret.

Algorithm 1 BARBAR: Bandit Algorithm with Robustness: Bad Arms get Recourse

- 1: **Parameters:** confidence $\delta \in (0, 1)$, time horizon T .
 - 2: Initialize $T_0 = 0$ and $\Delta_i^0 = 1$ for all $i \in [K]$.
 - 3: Set $\lambda = 1024 \ln(\frac{8K}{\delta} \log_2 T)$.
 - 4: **for** epochs $m = 1, 2, \dots$ **do**
 - 5: Set $n_i^m = \lambda(\Delta_i^{m-1})^{-2}$ for all $i \in [K]$.
 - 6: Set $N_m = \sum_{i=1}^K n_i^m$ and $T_m = T_{m-1} + N_m$.
 - 7: **for** $t = T_{m-1} + 1$ **to** T_m **do**
 - 8: choose an arm i with probability n_i^m / N_m and pull it.
 - 9: Set $r_i^m = S_i / n_i^m$ where S_i is the total reward from the pulls of arm i in this epoch.
 - 10: Set $r_*^m = \max_i \{r_i^m - \frac{1}{16} \Delta_i^{m-1}\}$
 - 11: Set $\Delta_i^m = \max\{2^{-m}, r_*^m - r_i^m\}$
-

Formally, the BARBAR (*Bandit Algorithm with Robustness: Bad Arms get Recourse*) algorithm is presented as [Algorithm 1](#). The following is our main result: an $\tilde{O}(KC + \sum_{i \neq i^*} 1/\Delta_i)$ regret bound for the BARBAR algorithm.

Theorem 1. *With probability at least $1 - \delta$, the regret of [Algorithm 1](#) is bounded by*

$$O\left(KC + \sum_{i \neq i^*} \frac{\log T}{\Delta_i} \log\left(\frac{K}{\delta} \log T\right)\right).$$

In [Appendix B](#) we give an example where our algorithm suffers $\Omega(KC)$ regret, showing that our analysis is essentially tight. To convert the bound of [Theorem 1](#) into a bound on the expected regret of the algorithm, one can set $\delta = 1/T$; then, with the remaining probability the regret is trivially bounded by T and the expected contribution to the regret is $O(1)$. Thus, the expected regret is bounded by

$$O\left(K \mathbb{E}[C] + \sum_{i \neq i^*} \frac{1}{\Delta_i} \log^2(KT)\right).$$

We now turn to prove [Theorem 1](#). We start with a simple observation.

Lemma 2. *The length N_m of epoch m satisfies*

$$\lambda 2^{2(m-1)} \leq N_m \leq K \lambda 2^{2(m-1)}.$$

Moreover, the number of epochs M is at most $\log_2 T$.

Proof: The arm i_m achieving the maximum in the definition of r_*^m has $r_*^m - r_{i_m}^m \leq 0$, so it has $\Delta_{i_m}^m = 2^{-m}$. This immediately implies that $N_{m+1} \geq n_{i_m}^{m+1} = \lambda(\Delta_{i_m}^m)^{-2} = \lambda 2^{2m}$. The upper bound follows from the fact that each Δ_i^m is at least 2^{-m} , and so $n_i^{m+1} \leq \lambda 2^{2m}$ for each i . The bound on the number of epochs M follows from the lower bound on N_m . \blacksquare

Let C_m^i be the random variable denoting the sum of the corruptions in epoch m to arm i 's reward, and let $C_m := \max_i C_m^i$. Let \tilde{n}_i^m be the random variable denoting the actual number of pulls of arm i in epoch m . We define an event \mathcal{E} that captures the relevant tail bounds as follows:

$$\mathcal{E} := \left\{ \forall m, i : |r_i^m - \mu_i| \leq \frac{2C_m}{N_m} + \frac{\Delta_i^{m-1}}{16} \text{ and } \tilde{n}_i^m \leq 2n_i^m \right\}. \quad (1)$$

Lemma 3. *It holds that $\Pr[\mathcal{E}] \geq 1 - \delta$.*

The lemma is obtained from the following concentration inequalities concerning the reward estimates r_i^m computed by the algorithm and the number of arm pulls \tilde{n}_i^m it uses.

Lemma 4. *For any fixed i, m and any $\beta \geq 4e^{-\lambda/16}$, we have*

$$\Pr \left[|r_i^m - \mu_i| \geq \frac{2C_m}{N_m} + \sqrt{\frac{4 \ln 4\beta}{n_i^m}} \right] \leq \beta \text{ and } \Pr[\tilde{n}_i^m \geq 2n_i^m] \leq \beta.$$

Proof: For this proof, we condition on all random variables before epoch m , so that n_i^m, N_m and T_{m-1} are deterministic quantities. At each step in this epoch, we pick arm i with probability $q_i^m := n_i^m / N_m$. Let Y_i^t be an indicator for arm i being pulled in step t . Let R_i^t be the stochastic reward of arm i on step t and $c_i^t := \tilde{R}_i^t - R_i^t$ be the corruption added to this arm by the adversary in this step. Note that c_i^t may depend on all the stochastic rewards up to (and including) round t , and also on all *previous* choices of the algorithm (though not the choice at step t). Finally, if we denote $E_m = [T_{m-1} + 1, \dots, T_m]$ be the N_m many time-steps in epoch m , then the first random variable we aim to control can be expressed as

$$r_i^m = \frac{1}{n_i^m} \sum_{t \in E_m} Y_i^t (R_i^t + c_i^t).$$

For ease of analysis, let us break the sum above into two, and define

$$A_i^m = \sum_{t \in E_m} Y_i^t R_i^t, \quad B_i^m = \sum_{t \in E_m} Y_i^t c_i^t.$$

Let us first bound the deviation of A_i^m . Observe that R_i^t is an independent draw from a $[0, 1]$ -valued r.v. with mean μ_i and Y_i^t is a independent draw from an $\{0, 1\}$ -valued r.v. with mean q_i^m . Moreover, we have that $\mathbb{E}[A_i^m] = N_m \cdot (q_i^m \mu_i) = n_i^m \mu_i \leq n_i^m$. Hence, by a Chernoff-Hoeffding bound (a multiplicative version thereof; see [Theorem 9](#) in [Appendix A](#)), we have

$$\Pr \left[\left| \frac{A_i^m}{n_i^m} - \mu_i \right| \geq \sqrt{\frac{3 \ln 4\beta}{n_i^m}} \right] \leq \frac{\beta}{2}. \quad (2)$$

Next we turn to bound the deviation of B_i^m . Consider the sequence of r.v.s X_1, \dots, X_T defined by $X_t = (Y_i^t - q_i^m) \cdot c_i^t$ for all t . Then $\{X_t\}_{t=1}^T$ is a martingale difference sequence with respect to the filtration $\{\mathcal{F}_t\}_{t=1}^T$ generated by the r.v.s $\{Y_j^s\}_{j \in [K], s \leq t}, \{R_j^s\}_{j \in [K], s \leq t+1}$. Indeed, since the corruption c_i^t becomes a deterministic value when conditioned on \mathcal{F}_{t-1} (as we assume a deterministic adversary), and since $\mathbb{E}[Y_i^t | \mathcal{F}_{t-1}] = q_i^m$, we have

$$\mathbb{E}[X_t | \mathcal{F}_{t-1}] = \mathbb{E}[Y_i^t - q_i^m | \mathcal{F}_{t-1}] \cdot c_i^t = 0.$$

Further, we have $|X_t| \leq 1$ for all t , and we can bound the predictable quadratic variation of this martingale as:

$$V = \sum_{t \in E_m} \mathbb{E}[X_t^2 \mid \mathcal{F}_{t-1}] \leq \sum_{t \in E_m} |c_i^t| \text{Var}(Y_t) \leq q_i^m \sum_{t \in E_m} |c_i^t|.$$

Applying a Freedman-type concentration inequality for martingales (see [Theorem 10](#) in [Appendix A](#)), we obtain that except with probability $\beta/4$,

$$\frac{B_i^m}{n_i^m} \leq \frac{q_i^m}{n_i^m} \sum_{t \in E_m} c_i^t + \frac{V + \ln 4\beta}{n_i^m} \leq 2 \frac{q_i^m}{n_i^m} \sum_{t \in E_m} |c_i^t| + \frac{\ln 4\beta}{n_i^m}.$$

Since $q_i^m = n_i^m / N_m$ and $\sum_{t \in E_m} |c_i^t| \leq C_m$, and further $n_i^m \geq \lambda \geq 16 \ln(4/\beta)$, we have with the same probability that

$$\frac{B_i^m}{n_i^m} \leq \frac{2C_m}{N_m} + \sqrt{\frac{\ln 4\beta}{16n_i^m}}.$$

Similar arguments show that $-B_i^m/n_i^m$ satisfies this bound except with probability $\beta/4$. It follows that

$$\Pr \left[\left| \frac{B_i^m}{n_i^m} \right| \geq \frac{2C_m}{N_m} + \sqrt{\frac{\ln 4\beta}{16n_i^m}} \right] \leq \frac{\beta}{2}. \quad (3)$$

We can now combine [Eqs. \(2\)](#) and [\(3\)](#) using a union bound to get the claimed bound, conditioned on the values of n_i^m, N_m, T_m . Finally, since the bound holds for any realizations of these r.v.s, it also holds unconditionally.

For the second claimed bound, we use a Chernoff-Hoeffding bound ([Theorem 9](#) in [Appendix A](#)) on the random variable $\tilde{n}_i^m = \sum_{t \in E_m} Y_t^i$, when we again condition on past epochs and later remove the conditioning. The expected value is $\mathbb{E}[\tilde{n}_i^m] = n_i^m \geq \lambda$, and the probability that the r.v. is more than twice its expectation is at most $2 \exp(-\tilde{n}_i^m/3) \leq 2 \exp(-\lambda/3) \leq \beta$. This completes the proof. ■

Proof of [Lemma 3](#): We use [Lemma 4](#) for each arm i and epoch m with $\beta = \delta / (2K \log_2 T)$; this value satisfies $\beta \geq 4e^{-\lambda/16}$. Since $n_i^m = \lambda(\Delta_i^{m-1})^{-2}$ and $\ln(4/\beta) = \lambda/1024 = \lambda/(4 \cdot 16^2)$, the lemma implies

$$\Pr \left[|r_i^m - \mu_i| \geq \frac{2C_m}{N_m} + \frac{\Delta_i^{m-1}}{16} \text{ and } \tilde{n}_i^m \geq 2n_i^m \right] \leq \frac{\delta}{K \log_2 T}.$$

A union bound over the K arms and the $\log_2 T$ epochs concludes the proof. ■

For the rest of the analysis, we will condition on the event \mathcal{E} . Our next step is to establish upper and lower bounds on our estimate of the maximal reward r_\star^m in epoch m .

Lemma 5. *Suppose that \mathcal{E} occurs. Then for all epochs m ,*

$$-\frac{2C_m}{N_m} - \frac{\Delta_{i_\star}^{m-1}}{8} \leq r_\star^m - \mu_\star \leq \frac{2C_m}{N_m}.$$

Proof: Recall that $r_\star^m = \max_i \{r_i^m - \frac{1}{16} \Delta_i^{m-1}\}$. In particular, $r_\star^m \geq r_{i_\star}^m - \frac{1}{16} \Delta_{i_\star}^{m-1}$; plugging in the lower bound on $r_{i_\star}^m$ implied by the event \mathcal{E} , we get the left inequality. On the other hand, plugging in the upper bound on r_i^m yields

$$r_\star^m \leq \max_i \left\{ \mu_i + \frac{2C_m}{N_m} + \frac{\Delta_i^{m-1}}{16} - \frac{\Delta_i^{m-1}}{16} \right\} = \max_i \{\mu_i\} + \frac{2C_m}{N_m} = \mu_\star + \frac{2C_m}{N_m}. \quad \blacksquare$$

We now establish upper and lower bounds on the gap estimate Δ_i^m of arm i on epoch m . For all epochs m , define the discounted corruption rate

$$\rho_m := \sum_{s=1}^m \frac{2C_s}{8^{m-s}N_s}. \quad (4)$$

Lemma 6. *Suppose that \mathcal{E} occurs. Then for all epochs m and arms i , it holds that*

$$\Delta_i^m \leq 2(\Delta_i + 2^{-m} + \rho_m).$$

Proof: The proof is by induction on m . For $m = 1$, the claim is trivially true for every i , because $\Delta_i^m \leq 1 \leq 2 \cdot 2^{-m}$. Next, suppose that the claim holds for $m - 1$. Using [Lemma 5](#) and the definition of the event \mathcal{E} (cf. [Eq. \(1\)](#)), we write

$$\begin{aligned} r_\star^m - r_i^m &= (r_\star^m - \mu^\star) + (\mu^\star - \mu_i) + (\mu_i - r_i^m) \\ &\leq \frac{2C_m}{N_m} + \Delta_i + \frac{2C_m}{N_m} + \frac{\Delta_i^{m-1}}{16}. \end{aligned}$$

Now using the induction hypothesis and the expression for ρ_{m-1} , we have

$$\begin{aligned} r_\star^m - r_i^m &\leq \Delta_i + \frac{4C_m}{N_m} + \frac{1}{8} \left(\Delta_i + 2^{-(m-1)} + \sum_{s=1}^{m-1} \frac{2C_s}{8^{m-1-s}N_s} \right) \\ &\leq 2 \left(\Delta_i + 2^{-m} + \sum_{s=1}^m \frac{2C_s}{8^{m-s}N_s} \right) \\ &= 2(\Delta_i + 2^{-m} + \rho_m). \end{aligned}$$

Since $\Delta_i^m = \max\{2^{-m}, r_\star^m - r_i^m\}$, the claim follows. ■

Armed with this result, we now prove a lower bound on Δ_i^m .

Lemma 7. *Suppose that \mathcal{E} occurs. Then for all epochs m and arms i , it holds that*

$$\Delta_i^m \geq \frac{1}{2}\Delta_i - 3\rho_m - \frac{3}{4}2^{-m}.$$

Proof: We write

$$\begin{aligned} \Delta_i^m &\geq r_\star^m - r_i^m \\ &\geq \left(\mu^\star - \frac{2C_m}{N_m} - \frac{\Delta_{i^\star}^{m-1}}{8} \right) - \left(\mu_i + \frac{2C_m}{N_m} + \frac{\Delta_i^{m-1}}{16} \right), \end{aligned}$$

where we use [Lemma 5](#) to give a lower bound on r_\star^m and [Eq. \(1\)](#) to give an upper bound on r_i^m . Since $\Delta_i = \mu^\star - \mu_i$, the above expression is at least

$$\geq \Delta_i - \frac{4C_m}{N_m} - \left(\frac{\Delta_{i^\star}^{m-1}}{8} + \frac{\Delta_i^{m-1}}{16} \right),$$

which using [Lemma 6](#) can be further lower bounded by

$$\begin{aligned} &\geq \Delta_i - \frac{4C_m}{N_m} - \left(\frac{3}{8}\rho_{m-1} + \frac{3}{8}2^{-(m-1)} + \frac{1}{8}\Delta_i \right) \\ &\geq \frac{\Delta_i}{2} - 3 \sum_{s=1}^m \frac{2C_s}{8^{m-s}N_s} - \frac{3}{4}2^{-m}. \end{aligned}$$

Recalling the expression for ρ_m , the lemma follows. \blacksquare

We are now ready to bound the regret of [Algorithm 1](#) and prove our main result.

Proof of Theorem 1: We will prove that the claimed bound holds true under the event \mathcal{E} , which occurs with probability at least $1 - \delta$.

We decompose the total regret across epochs and within epochs, to the arms pulled. There are \tilde{n}_i^m pulls of arm i in epoch m , each causing a regret of $(\mu^* - \mu_i)$. Thus the total regret can be written as

$$\sum_{m=1}^M \sum_{i=1}^K (\mu^* - \mu_i) \tilde{n}_i^m \leq 2 \sum_{m=1}^M \sum_{i=1}^K (\mu^* - \mu_i) n_i^m = 2 \sum_{m=1}^M \sum_{i=1}^K \Delta_i n_i^m. \quad (5)$$

Fix an epoch m and an arm i and denote $\mathcal{R}_i^m := \Delta_i n_i^m$. Recall that $n_i^m = \lambda(\Delta_i^{m-1})^{-2}$, and consider three cases as follows.

Case 1: $0 < \Delta_i \leq 4/2^m$. In this case we use the upper bound $n_i^m \leq \lambda 2^{2(m-1)}$, which follows from $\Delta_i^{m-1} \geq 2^{m-1}$. This yields

$$\mathcal{R}_i^m \leq \frac{4\lambda}{\Delta_i^2} \Delta_i = \frac{4\lambda}{\Delta_i}. \quad (6)$$

Case 2: $\Delta_i > 4/2^m$ and $\rho_{m-1} < \Delta_i/32$. In this case, [Lemma 7](#) implies that

$$\Delta_i^{m-1} \geq \frac{1}{2}\Delta_i - 3\rho_{m-1} - \frac{3}{4}2^{-(m-1)} \geq \Delta_i \left(\frac{1}{2} - \frac{3}{32} - \frac{3}{8} \right) = \frac{1}{32}\Delta_i.$$

In turn, we have $n_i^m = \lambda(\Delta_i^{m-1})^{-2} \leq 32^2\lambda/\Delta_i^2$ from which it follows that

$$\mathcal{R}_i^m \leq \frac{32^2\lambda}{\Delta_i}. \quad (7)$$

Case 3: $\rho_{m-1} \geq \Delta_i/32$ (and $\Delta_i > 4/2^m$). This rearranges to $\Delta_i \leq 32\rho_{m-1}$. We use again the upper bound $n_i^m \leq \lambda 2^{2(m-1)}$ so that

$$\mathcal{R}_i^m \leq \lambda \Delta_i 2^{2(m-1)} \leq 8\lambda \rho_{m-1} 2^{2m}. \quad (8)$$

Considering the three cases, from [Eqs. \(6\) to \(8\)](#) we have for all arms $i \neq i^*$ and epochs m that

$$\mathcal{R}_i^m \leq \frac{32^2\lambda}{\Delta_i} + 8\lambda \rho_{m-1} 2^{2m},$$

and summing this over all epochs and arms, we have that the sum $\sum_{m=1}^M \sum_{i=1}^K \mathcal{R}_i^m$ is bounded by

$$32^2 \lambda \sum_{i \neq i^*} \frac{\log_2 T}{\Delta_i} + 8\lambda \sum_{i \neq i^*} \left(\sum_{m=1}^M \rho_{m-1} 2^{2m} \right).$$

For bounding the second term above, observe that

$$\begin{aligned} \sum_{m=1}^M \rho_{m-1} 2^{2m} &\leq \sum_{m=1}^M 2^{2m} \left(\sum_{s=1}^{m-1} \frac{2C_s}{8^{m-1-s} N_s} \right) \\ &= 2 \sum_{s=1}^M C_s \sum_{m=s}^M \frac{2^{2m}}{8^{m-1-s} N_s}; \end{aligned}$$

Since $N_s \geq \lambda 2^{2(s-1)}$ (recall [Lemma 2](#)), the inner sum is upper bounded by

$$\frac{16}{\lambda} \sum_{m=s}^M \frac{4^{m-1-s}}{8^{m-1-s}} \leq \frac{16}{\lambda} \sum_{j=1}^{\infty} 2^{-j} \leq \frac{16}{\lambda}.$$

Combining the inequalities and plugging in the value of λ gives the claimed regret bound. \blacksquare

4. Better Bounds for Special Cases

In this section, we study a few special cases of the main model and show that under various kinds of additional assumptions, we can obtain an improved regret bound in which the first term $O(KC)$ is improved to $\tilde{O}(C)$. Some of these models are interesting in their own right; others serve to suggest that a better bound of the form $\tilde{O}(C + \sum_{i \neq i^*} 1/\Delta_i)$ may be achievable in the general case.

4.1. Corruption at a fixed unknown rate

Consider the scenario each round is contaminated independently and uniformly at random by the adversary with probability $\eta \in (0, 1)$, which may be unknown to the player. From the point of view of a stochastic bandits algorithm (such as Active Arm Elimination), this case is equivalent to operating over an instance where the means μ_i 's are adversarially perturbed by $\pm\eta$. (This again can be seen from [Lemma 4](#) with $C_m = \eta N_m$.) A regret bound of $\tilde{O}(\eta T + \sum_{i \neq i^*} 1/\Delta_i)$ then follows.

4.2. Known corruption level

Now suppose that the corruption amount C is known to the player. In this case, consider the classical Active Arm Elimination algorithm, with the modification that epoch m has length at least $N_m = \Omega(C + 2^m \log(K/\delta))$ which is ensured by pulling the winning arm sufficiently many times. Even if all C corruptions fall in epoch m , [Lemma 4](#) shows that the algorithm's empirical estimates of the means are accurate to within $\tilde{O}(2^{-m})$. It can be then verified that an arm with gap $\Delta_i \geq 2^{-m}$ will be eliminated before epoch $m + O(1)$ with high probability, leading to a bound of $\tilde{O}(C + \sum_{i \neq i^*} 1/\Delta_i)$.

4.3. Corruption on an unknown prefix

Assume that the adversary's corruption is known to take place only in the first C steps, where C may be unknown. In this case, the simple algorithm that plays AAE for steps $[2^m, 2^{m+1})$, starting afresh each time has regret $\tilde{O}(C + \sum_{i \neq i^*} 1/\Delta_i)$.

4.4. Known minimal gap

We can also obtain a similarly improved bound when the minimal gap $\Delta = \min_{i \neq i^*} \Delta_i$ is known. Consider the following variant of our algorithm: in each epoch m , play last epoch's winning arm 2^{2m} times and play every other arm λ/Δ^2 times. This algorithm gets regret K/Δ in each epoch from all arms except the winning one. If the winning arm is i^* , it causes no regret in this epoch. For the winner to be an arm $i \neq i^*$, the corruption in epoch $m-1$ must be at least $C_m \geq (\Delta_i - \Delta/8)2^{2(m-1)} \geq \Delta_i 2^{2(m-2)}$. This results in an additional regret of $O(\Delta_i 2^{2m})$. Thus, the overall regret is $\tilde{O}(C + K\lambda/\Delta)$ for any corruption level C .

4.5. Known maximal reward

Finally, consider the case where the player knows the value of μ^* . For this case, we change the algorithm so that $\Delta_i^m = \max\{2^{-m}, \mu^* - r_i^m, \Delta_i^{m-1}/2\}$; i.e., we use our knowledge of the optimal bias, and moreover our estimates of the gaps does not drop very fast. In this case, we prove the following variant of [Lemma 7](#):

Lemma 8. *Suppose that \mathcal{E} occurs. Then for all epochs m and arms i , it holds that*

$$\Delta_i^m \geq \frac{\Delta_i}{2} - \frac{2C_m}{N_m}.$$

Proof: If $\Delta_i^{m-1} \geq \Delta_i$ then the change in the algorithm ensures that $\Delta_i^{m-1} \geq \Delta_i^m/2$. Otherwise, [Eq. \(1\)](#) ensures that $|r_i^m - \mu_i| \leq 2C_m/N_m + \Delta_i/16$. Now

$$\Delta_i^m \geq \mu^* - r_i^m \geq (\mu^* - \mu_i) - \frac{2C_m}{N_m} - \frac{\Delta_i}{16} \geq \frac{\Delta_i}{2} - \frac{2C_m}{N_m}.$$

Hence the claim follows. ■

To bound the regret, the proof is similar to [Theorem 1](#). Consider an arm i and epoch m such that $C_m/N_m \leq \Delta_i/8$. [Lemma 8](#) implies that $\Delta_i^m \geq \Delta_i/4$, and hence the regret due to this arm i in the next epoch $m+1$ is $n_i^{m+1}\Delta_i = \lambda(\Delta_i^m)^{-2}\Delta_i \leq O(\lambda/\Delta_i)$. Summing this over all arms and epochs gives us $\tilde{O}(\sum_{i \neq i^*} 1/\Delta_i)$. Next, for arms i and epochs m where $C_m/N_m \geq \Delta_i/8$, the regret in the next round is $n_i^{m+1}\Delta_i \leq 4n_i^m\Delta_i$, since we ensured that $\Delta_i^m \geq \Delta_i^{m-1}/2$. Using the fact that $\Delta_i \leq 8C_m/N_m$, we have the regret is at most $32n_i^m C_m/N_m$. Summing over all the arms in this epoch $m+1$ gives us $32C_m$, and then summing over all rounds gives us $32C$. Hence, if we know μ^* , the total regret is at most $O(C) + \tilde{O}(\sum_{i \neq i^*} 1/\Delta_i)$.

5. Discussion

We conclude with a discussion of various aspects of our model and results.

5.1. A Lower bound

We show an $\Omega(C)$ lower bound for the pseudo-regret of any algorithm. Consider the case of two arms with deterministic rewards 0 and 1, and consider an adversary that at each time step swaps the rewards of the arms with probability $1/2$. The expected total amount of corruption is therefore $T/2$, and it makes the arms appear indistinguishable to the algorithm. Thus, with constant probability the algorithm will suffer pseudo-regret of $\Omega(C)$. This lower bound for $C = T/2$ can be extended to smaller C by placing this instance in the first $2C$ steps. In other words, for the pseudo-regret under this contamination model there is a lower bound of $\Omega(C)$ that holds for all C up to $T/2$. Moreover this lower bound holds even for the case of known C .

5.2. Alternative notions of regret

So far we have focused on the pseudo-regret metric, which is the most common in stochastic bandit problems. One might consider several other possible notions of regret. First, the realized regret with respect to the actual corrupted rewards may be of interest:

$$\mathcal{R}'(T) = \max_{i \in [K]} \sum_{t=1}^T (\tilde{R}_i^t - \tilde{R}_{i_t}^t).$$

Under this metric, not only do the adversarial corruptions perturb the feedback observed by the player, but they also affect the rewards she gains. As we discuss in [Appendix C](#), any high-probability bound on the pseudo-regret also implies a high-probability bound on the realized regret, and our upper bounds transfer directly (up to a logarithmic factor) to this notion as well.

Note that for this notion of regret one could also attain an $O(\sqrt{T})$ -type regret bound, at any level of corruption C , using an adversarial bandits algorithm like EXP3. On the other hand, [Lykouris et al. \(2018\)](#) show a lower bound of $\Omega(C^{1-\delta})$ on the regret for any algorithm that gets regret $O(\log^{2-\varepsilon}(T)/\Delta)$ for the uncorrupted case, for any $\varepsilon, \delta > 0$. This lower bound still leaves open the possibility of obtaining a “best of all worlds” type guarantee that smoothly interpolates between the favorable stochastic guarantees and the robust adversarial results, which is meaningful even for $C = \Omega(\sqrt{T})$. (For example, one might allow for slightly suboptimal $O(\text{polylog}(T)/\Delta)$ regret in the non-corrupted case, which invalidates the latter lower bound.)

However, recalling the original motivation for our problem, the above realized regret is arguably not the “correct” notion. In the motivating scenarios, such as click fraud or recommendation poisoning by fake reviewers, the rewards coming from the corruptions (the *outliers*) are not of interest, and in fact, the goal is to perform well on the non-corrupted rewards (the *inliers*). The corresponding definition of regret for such cases would take into account only the non-corrupted steps:

$$\mathcal{R}'_{\text{in}}(T) = \max_{i \in [K]} \sum_{t \notin \mathcal{C}} (\tilde{R}_i^t - \tilde{R}_{i_t}^t),$$

where \mathcal{C} is the set of steps in which there was some corruption.² Our upper bounds, as well as the simple $\Omega(C)$ lower bound argued above, immediately extend to this notion as well.

As a consequence, this definition of regret does not admit $o(C)$ -regret bounds for any C , and one cannot hope for a robust $O(\sqrt{T})$ -type regret guarantee that would hold for any level of corruption. How do we square this with the $O(\sqrt{T})$ -type upper bounds that hold for the classical definition of regret, achieved by adversarial algorithms like EXP3? On the instance used to show the $\Omega(C)$ lower bound above, those algorithms will balance the linear regret on the inlier steps with a linear negative regret on the outlier steps.

5.3. Open questions

Our work leaves a gap between the lower bound of $\Omega(C)$ and the upper bound of $O(KC)$, closing which is a compelling open question. Improving the $\log^2 T$ dependence in our expected regret bound to $\log T$ is another interesting problem for future investigation. Finally, it would also be interesting to extend our results to other online prediction settings, e.g., linear bandits and contextual bandits.

² Here we assume a slightly more restrictive definition of the corruption level, where C is the total number of rounds corrupted by the adversary. This definition is consistent with most of the motivating scenarios discussed.

Acknowledgments

AG is supported in part by NSF awards CCF-1536002, CCF-1540541, and CCF-1617790.

References

- Jason Altschuler, Victor-Emmanuel Brunel, and Alan Malek. Best arm identification for contaminated bandits. *arXiv preprint arXiv:1802.09514*, 2018.
- Peter Auer and Chao-Kai Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 116–120, 2016.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002a.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multi-armed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002b.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15, pages 19–26, 2011.
- Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 42–1, 2012.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, and Gábor Lugosi. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013.
- Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 47–60. ACM, 2017.
- Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 655–664. IEEE, 2016.
- Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Learning geometric concepts with nasty noise. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1061–1073. ACM, 2018.
- Devdatt P Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(Jun):1079–1105, 2006.
- Peter J. Huber. Robust estimation of a location parameter. *Ann. Math. Statist.*, 35(1):73–101, 03 1964.

- Sayash Kapoor, Kumar Kshitij Patel, and Purushottam Kar. Corruption-tolerant bandit learning. *Machine Learning*, pages 1–29, 2018.
- Michael Kearns and Ming Li. Learning in the presence of malicious errors. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 267–280, New York, NY, USA, 1988. ACM.
- Adam Klivans, Pravesh K Kothari, and Raghu Meka. Efficient algorithms for outlier-robust regression. In *COLT*, 2018.
- Kevin A Lai, Anup B Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 665–674. IEEE, 2016.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 114–122, New York, NY, USA, 2018. ACM.
- Yevgeny Seldin and Gábor Lugosi. An improved parametrization and analysis of the EXP3++ algorithm for stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 1743–1759, 2017.
- Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *ICML*, pages 1287–1295, 2014.
- Han Shao, Xiaotian Yu, Irwin King, and Michael R Lyu. Almost optimal algorithms for linear stochastic bandits with heavy-tailed payoffs. In *Advances in Neural Information Processing Systems*, pages 8430–8439, 2018.
- L.G. Valiant. Learning disjunctions of conjunctions. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 560–566, 1985.
- Xiaotian Yu, Han Shao, Michael R Lyu, and Irwin King. Pure exploration of multi-armed bandits with heavy-tailed payoffs. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 937–946, 2018.
- Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. *arXiv preprint arXiv:1807.07623*, 2018.
- Julian Zimmert, Haipeng Luo, and Chen-Yu Wei. Beating stochastic and adversarial semi-bandits optimally and simultaneously. *arXiv preprint arXiv:1901.08779*, 2019.

Appendix A. Concentration Inequalities

In our analysis we require multiplicative versions of standard concentration inequalities, that provide sharper rates for random variables with small expectations. First, we state a standard multiplicative variant of the Chernoff-Hoeffding bound.

Theorem 9 (e.g., [Dubhashi and Panconesi 2009, Theorem 1.1](#)). *Suppose that X_1, \dots, X_T are independent $[0, 1]$ -valued random variables, and let $X = \sum_{t=1}^T X_t$. Then for any $\varepsilon > 0$,*

$$\Pr[|X - \mathbb{E}[X]| \geq \varepsilon \mathbb{E}[X]] \leq 2 \exp\left(-\frac{\varepsilon^2}{3} \mathbb{E}[X]\right).$$

Equivalently, for any $\delta > 0$,

$$\Pr\left[|X - \mathbb{E}[X]| \leq \sqrt{3\mathbb{E}[X] \ln \frac{2}{\delta}}\right] \geq 1 - \delta.$$

Next, we record a variant of Freedman's martingale concentration inequality, taken from [Beygelzimer et al. \(2011\)](#). (Technically, the concentration inequality in [Beygelzimer et al. \(2011\)](#) is stated for the filtration generated by the martingale itself; the proof however extends straightforwardly to more general filtrations.)

Theorem 10 ([Beygelzimer et al. 2011, Theorem 1](#)). *Suppose that X_1, \dots, X_T is a martingale difference sequence with respect to a filtration $\{\mathcal{F}_t\}_{t=1}^T$, and let $X = \sum_{t=1}^T X_t$. Assume that $|X_t| \leq b$ for all t , and define $V = \sum_{t=1}^T \mathbb{E}[X_t^2 \mid \mathcal{F}_{t-1}]$. Then for any $\delta > 0$,*

$$\Pr\left[X \leq \frac{V}{b} + b \ln \frac{1}{\delta}\right] \geq 1 - \delta.$$

For completeness, we give a short proof of a slightly weaker bound.

Proof: We will prove a slightly weaker bound with a deviation of $V/b + 4b \ln(2/\delta)$. Freedman's inequality says that for any $\lambda \geq 0, \sigma^2 > 0$,

$$\Pr[X \geq \lambda \wedge V \leq \sigma^2] \leq \exp\left(-\frac{\lambda^2/2}{\sigma^2 + b\lambda/3}\right).$$

Setting $\sigma^2 = b\lambda$, we have for all $\lambda > 0$ that $\Pr[X \geq \lambda \geq V/b] \leq e^{-\lambda/4b}$. We now apply a discretization argument over λ . Fix $\varepsilon > 0$, and note that if $X \geq V/b + \varepsilon$ then there is an integer $n \geq 1$ such that $X \geq n\varepsilon \geq V/b$. Thus, by a union bound we have for any $\varepsilon \geq 4b \ln 2$:

$$\Pr\left[X \geq \frac{V}{b} + \varepsilon\right] \leq \sum_{n=1}^{\infty} \Pr\left[X \geq n\varepsilon \geq \frac{V}{b}\right] \leq \sum_{n=1}^{\infty} e^{-n\varepsilon/4b} = \frac{e^{-\varepsilon/4b}}{1 - e^{-\varepsilon/4b}} \leq 2e^{-\varepsilon/4b}.$$

Choosing $\varepsilon = 4b \ln(2/\delta)$ makes the right-hand side less than δ . ■

Appendix B. Tightness of Our Analysis

Here we give an example where our algorithm suffers regret $\Omega(CK)$, showing that the first term in our analysis is tight. Consider the case where the rewards of the arms are $\mu_1 = 1$ and $\mu_i = 0$ for all $i \neq 1$. In the uncorrupted case, for each epoch $m \geq 1$ we have $\Delta_1^m = 2^{-m}$ and $\Delta_i^m = r_\star^m = 1 - \frac{1-(1/16)^m}{15} \in [14/15, 15/16]$. Consequently, the m^{th} epoch has length $N_m = \lambda(2^{2(m-1)} + O(K))$ as long as we see no corruptions.

Suppose $C = N_c = \lambda(2^{2(c-1)} + O(K))$ for some $c > 2 \log_2 K$, say. In our example, all the corruptions happen in epoch c . Specifically, we corrupt the high-value arm to also have payoff zero in all N_c time-steps in this epoch. This causes $r_\star^c \leq 0$ and all $\Delta_i^c = 2^{-c}$: in essence we lose all information from the previous epochs. Therefore the length of the next epoch becomes $N_{c+1} = \lambda K 2^{2c} \approx 2KC$ which is K times as much; moreover, each arm is pulled $\lambda 2^{2c} \approx C$ times in this epoch. The regret just from this epoch is $\approx (K-1)C$, which shows we indeed need the factor of K multiplying the C term in the regret for our algorithm.

Appendix C. From Pseudo-regret to Regret

Here we show how a (high probability) bound over the pseudo regret can be converted to a bound into a (high probability) bound over the actual, realized regret in a black-box fashion. We assume the regret is defined based on the uncorrupted stochastic rewards; the analogous notion that uses the corrupted rewards is only an additive $O(C)$ larger in the worst case. As a consequence, our (positive) results regarding the pseudo regret immediately extend to the realized regret with respect to either the corrupted or uncorrupted rewards.

Using the notation established in [Section 2](#), recall that the pseudo regret is given by

$$\mathcal{R} = \sum_{i=1}^T (\mu^\star - \mu_i) = \sum_{i \neq i^\star} T_i \Delta_i,$$

where T_i is the number of times arm i was pulled. On the other hand, the (realized) regret is the quantity

$$\mathcal{R}' = \max_{i \in [K]} \sum_{t=1}^T (R_i^t - R_{i^\star}^t).$$

Suppose that we have established a bound $\Pr(\mathcal{R} \leq B) \geq 1 - \delta$ on the pseudo regret, where $B \geq \max_{i \neq i^\star} 1/\Delta_i$. We want to establish a high-probability bound on \mathcal{R}' . Define the quantity:

$$\mathcal{R}^\star = \sum_{t=1}^T (R_{i^\star}^t - R_{i^\star}^t).$$

We will separately bound $\mathcal{R}' - \mathcal{R}^\star$ and $\mathcal{R}^\star - \mathcal{R}$. Applying Azuma's inequality on the random variable $\mathcal{R}^\star - \mathcal{R}$ we have with probability at least $1 - \delta$ that

$$\mathcal{R}^\star - \mathcal{R} \leq \sqrt{\sum_{i \neq i^\star} T_i \ln \frac{1}{\delta}}.$$

Applying Holder's inequality,

$$\mathcal{R}^\star - \mathcal{R} \leq \sqrt{\sum_{i \neq i^\star} T_i \Delta_i \cdot \frac{1}{\Delta_i} \ln \frac{1}{\delta}} = \sqrt{\left(\sum_{i \neq i^\star} T_i \Delta_i \right) \cdot \left(\max_{i \neq i^\star} \frac{1}{\Delta_i} \ln \frac{1}{\delta} \right)}.$$

Now, since B is a bound over the pseudo-regret we have that $\mathcal{R} = \sum_{i \neq i^*} T_i \Delta_i \leq B$ with probability at least $1 - \delta$. We have therefore shown that with probability at least $1 - 2\delta$,

$$\mathcal{R}^* - \mathcal{R} \leq \sqrt{B \max_{i \neq i^*} \frac{1}{\Delta_i} \ln \frac{1}{\delta}}.$$

On the other hand, another application of Azuma's inequality shows that with probability $\geq 1 - \delta$,

$$\mathcal{R}' - \mathcal{R}^* \leq \max_{i \in [K]} \left(-\Delta_i T + \sqrt{T \ln \frac{K}{\delta}} \right).$$

If $\min_{i \neq i^*} \Delta_i \geq \sqrt{\ln(K/\delta)/T}$, then this difference is negative. Otherwise, it is bounded as

$$\mathcal{R}' - \mathcal{R}^* \leq \sqrt{T \ln \frac{K}{\delta}} \leq \max_{i \neq i^*} \frac{1}{\Delta_i} \ln \frac{K}{\delta}.$$

Also, we have $\max_{i \neq i^*} 1/\Delta_i \leq B$ by assumption. It follows that with probability at least $1 - 3\delta$,

$$\mathcal{R}' \leq B \left(1 + \sqrt{\ln \frac{1}{\delta} + \ln \frac{K}{\delta}} \right).$$

Note also that if $B^{-1} \ln \frac{K}{\delta} \cdot (\max_{i \neq i^*} 1/\Delta_i)$ is bounded by a constant (as it typically is), then \mathcal{R}' is at most $O(B)$.