Stochastic Bandits Robust to Adversarial Corruptions

Thodoris Lykouris

Cornell University, USA

Vahab Mirrokni

Google Research, USA

Renato Paes Leme

Google Research, USA

ABSTRACT

We introduce a new model of stochastic bandits with adversarial corruptions which aims to capture settings where most of the input follows a stochastic pattern but some fraction of it can be adversarially changed to trick the algorithm, e.g., click fraud, fake reviews and email spam. The goal of this model is to encourage the design of bandit algorithms that (i) work well in mixed adversarial and stochastic models, and (ii) whose performance deteriorates gracefully as we move from fully stochastic to fully adversarial models.

In our model, the rewards for all arms are initially drawn from a distribution and are then altered by an adaptive adversary. We provide a simple algorithm whose performance gracefully degrades with the total corruption the adversary injected in the data, measured by the sum across rounds of the biggest alteration the adversary made in the data in that round; this total corruption is denoted by C . Our algorithm provides a guarantee that retains the optimal guarantee (up to a logarithmic term) if the input is stochastic and whose performance degrades linearly to the amount of corruption C , while crucially being agnostic to it. We also provide a lower bound showing that this linear degradation is necessary if the algorithm achieves optimal performance in the stochastic setting (the lower bound works even for a known amount of corruption, a special case in which our algorithm achieves optimal performance without the extra algorithm).

CCS CONCEPTS

• **Theory of computation** → **Online learning algorithms; Online learning theory; Regret bounds; Adversary models;**

KEYWORDS

Bandits, Online learning, Corruptions, Click fraud

ACM Reference Format:

Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. 2018. Stochastic Bandits Robust to Adversarial Corruptions. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC'18)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3188745.3188918>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

STOC'18, June 25–29, 2018, Los Angeles, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5559-9/18/06.

<https://doi.org/10.1145/3188745.3188918>

1 INTRODUCTION

In online learning with bandit feedback, a learner needs to decide at each time between alternative actions or *arms* of unknown quality, facing a trade-off between *exploiting* profitable past actions or *exploring* new actions about which she has little information. Bandit problems are typically classified according to how the rewards are generated. In stochastic bandits, rewards are drawn from fixed but unknown distributions, which models settings where the alternatives follow particular patterns and do not react to the learner. The other extreme is adversarial bandits, which are robust to rewards that are specifically designed to trick the learner, as in game-theoretic settings.

In this paper, we focus on settings where the overall behavior is essentially stochastic but a small fraction of the rewards can be adversarially changed. Classic stochastic bandit algorithms, like Upper Confidence Bound (UCB) [ACBF02] or Active Arm Elimination (AAE) [EMM06], base most of their decisions on a few observations made in an initial phase of the algorithm and therefore can be easily tricked into incurring linear regret if very few arms are corrupted. Adversarial bandit algorithms like EXP3 are not fooled by such tricks, but cannot exploit the fact that the input is mostly stochastic.

Our goal is to robustify the stochastic setting by designing algorithms that can tolerate corruptions and still be able to exploit the stochastic nature of the input. The algorithms we design are *agnostic* to the corruption, i.e. they can tolerate any level of corruption, and the guarantee degrades gracefully as more corruption is added. Moreover, we prove lower bounds showing that our results are tight up to a logarithmic factor. Before we explain our technical contribution in detail, we describe examples of settings we have in mind.

Click fraud. In pay-per-click online advertising, the platform selects for each pageview an ad to display and obtains a certain reward if the user clicks on the ad. The click probabilities are unknown. The tension between repeatedly displaying a particular profitable ad that provides reliable revenue and exploring other potentially more rewarding options is a major application of stochastic bandits in the ads industry.

If it weren't for a phenomenon known as *click fraud*, this would be a textbook example of stochastic bandits. In click fraud, botnets maliciously simulate users clicking on an ad to trick learning algorithms. One example is a bot consistently making searches to trigger some ad and not clicking on it to make it seem like a certain ad has very low click-through-rate in order to boost its competitor.

Recommendation systems. A platform recommending activities or services to a user faces the same trade-off. Suggesting new restaurants leads to faster learning of the best spots but may result to

dissatisfaction of the customers who are led to disappointing experiences. While most inputs follow a stochastic pattern, some inputs are typically corrupted: either maliciously, e.g. fake reviews by competitors, or non-maliciously, e.g. construction next-door makes the restaurant less desirable in certain interval. This corruption may again exhibit arbitrary patterns and is not identically distributed over time, yet it is dwarfed by the fact that most of the input is stochastic.

There are several other such examples: emails mostly follow a stochastic pattern except a fraction of them which are spam and are designed to trick algorithms. Internet searches follow a predictable pattern except certain spikes caused by unpredictable events. Data collection used in the econometric process often suffers from errors that affect a small part of the input. In all those cases, the vast majority of the input follows a predictable pattern, but a fraction of the samples are corrupted.

1.1 Our Contribution

Our model. In this paper, we introduce a new model of stochastic bandits with adversarial corruptions. The goal of this model is to encourage the design of bandit algorithms (i) that work well in mixed adversarial and stochastic models, and (ii) whose performance deteriorates gracefully as we move from fully stochastic to fully adversarial models.

In this model there are K arms, each associated with a fixed reward distribution $\mathcal{F}(a)$. At each round t , a random reward $r_S^t(a) \sim \mathcal{F}(a)$ is drawn and an adversary can change the reward to $r^t(a)$, possibly using information about the realizations of $r_S^\tau(a)$ from both the current and previous rounds $\tau \leq t$ as well as the probability that the learner puts on each arm. The learner then draws an arm a^t and obtains $r^t(a^t)$ both as reward and feedback. We say that the adversary is C -corrupted if in every sample path we have $\sum_t \max_a |r^t(a) - r_S^t(a)| \leq C$.

Our results. The main result (Theorem 3 in Section 3) is a learning algorithm we term Multi-layer Active Arm Elimination Race that with probability $1 - \delta$ has regret

$$O\left(\sum_{a \neq a^*} \frac{K \cdot C \log(KT/\delta) + \log(T)}{\Delta(a)} \cdot \log(KT/\delta)\right)$$

where $\Delta(a)$ is the *gap* of arm a , i.e. the difference in stochastic means of arm a and the optimal arm a^* . For arms with very small gap, i.e. when $\Delta(a) \leq 1/\sqrt{T}$, the inverse dependence on the gap can be replaced by $\sqrt{T}$. It is possible to improve the bound by a log factor for pseudo-regret, i.e. maximum expected regret against any fixed arm, obtaining: $O\left(\sum_{a \neq a^*} \frac{K \cdot C + \log(T)}{\Delta(a)} \cdot \log(KT)\right)$. Two important features of the algorithm are that the guarantee is:

- **Agnostic:** The algorithm does not need to know the corruption level C . The guarantee is provided with respect to how much corruption was added in retrospect. If the corruption level is known, we can remove the dependence on $K \cdot \log(KT/\delta)$ as shown in Theorem 1.

- **High Probability:** Our bounds hold with high probability which is important for practical applications as the ones described above. In contrast, the weaker definition of pseudo-regret often hides events with large regret that are offset by events with large negative regret.

The stochastic case corresponds to $C = 0$ in which case we recover a bound that is slightly worse than the guarantee provided by UCB. Our algorithm obtains $O\left(\sum_{a \neq a^*} \log(T) \cdot \log(T/\delta)/\Delta(a)\right)$ with probability $1 - \delta$, while UCB obtains this bound without the $\log(T)$ term.

En route to the result, in Theorem 2 we show an algorithm that, for any fixed known C , provides regret $O\left(\sum_{a \neq a^*} \frac{\log(KT/\delta)}{\Delta(a)}\right)$ for stochastic input and $O\left(K \cdot C \cdot \sum_{a \neq a^*} \frac{(\log(KT/\delta))^2}{\Delta(a)}\right)$ if it is C -corrupted. In other words, if we only need to tolerate either a known level C or zero corruptions, we save a logarithmic factor from the bound, and match the bound provided by UCB in the stochastic case.

Another question is whether the linear dependence on the corruption level is tight. In Section 4, we show that it cannot be improved upon without decay in the stochastic guarantee (i.e. while still guaranteeing logarithmic regret when the input is stochastic). The lower bound is an adaptation from the adversarial to the corrupted setting of a result from Auer and Chiang [AC16]. This holds even for the case where the corruptions are either 0 or a known level C (where our algorithm provides a matching upper bound). We prove in Theorem 4 that an algorithm with pseudo-regret $O(\log(T)/\Delta)$ in the stochastic setting ($C = 0$) then for every constant $\epsilon > 0$, there is a $O(T^\epsilon)$ -corrupted instance where the algorithm incurs regret $\Omega(T^\epsilon)$ with constant probability.

Our algorithm can also be viewed through the lens of the best of both worlds literature [BS12, SS14, AC16, SL17], where the goal is to design algorithms that simultaneously provide logarithmic regret guarantees in the stochastic regime and square-root guarantees in the adversarial. In Section 5, we sketch how our algorithm can be appropriately modified to obtain, for any constant $0 < a < 1/2$, $\tilde{O}(C)$ pseudo-regret for $C = O(T^a)$ and $\tilde{O}(T^{a+1/2})$ pseudo-regret otherwise. We observe that the results in the best of both worlds literature correspond to the case where $a = 0$. We note that such bounds are obtained for pseudo-regret and not regret with high-probability.

Our techniques. The starting point of our design are classical stochastic bandit learning algorithms like UCB and Active Arm Elimination. Such algorithms are very susceptible to corruptions since they base most of their decisions on a small initial exploration phase. Therefore, with a small number of corruptions it is possible to completely trick the algorithm into eliminating the optimal arm.

We address this issue by robustifying them using a multi-layer approach. The learning algorithm consists of multiple layers running in parallel. The layers have decreasing speed and increasing tolerance to corruption. The first layer finishes very fast selecting an arm as optimal, but provides no tolerance to corruption. Subsequent layers are more robust but also slower.

The resulting algorithm is a *race* between different layers for picking the optimal arm. Once the fastest layer finishes, it provides a first crude estimate of the optimal arm. Once slower layers finish, we obtain finer and finer estimates of the optimal arm.

Our second main idea is that we can obtain more robust algorithms by *subsampling*. If a layer is only selected with probability p , it only receives in expectation a p -fraction of the corruption injected by the adversary. If p is low enough, the layer behaves almost as if it was stochastic.

Finally, we couple the different layers together by a process of *global eliminations*. This process enables slower layers to eliminate arms in faster layers. Such a process is necessary for preventing inaccurate layers from pulling suboptimal arms too often.

1.2 Related Work

Online learning with stochastic rewards goes back to the seminal work of Lai and Robbins [LR85]. The case of adversarial rewards was introduced by Auer et al. [ACBFS03]. The reader is referred to the books of Cesa-Bianchi and Lugosi [CBL06], Bubeck and Cesa-Bianchi [BCB12], and Slivkins [Sl17] for an elaborate overview of the area. These two extremes suffer from orthogonal problems; the one is overoptimistic expecting that all rewards come from the same distribution while the other one is too pessimistic in order to be protected against malicious adversaries. Our work addresses the middle ground: rewards come from distributions but are often adversarially corrupted. This is motivated by the non-robustness of stochastic learning algorithms to even small corruption levels.

Closely related to our work lie the works on best of both worlds guarantees [BS12, SS14, AC16, SL17]. These works achieve (up to logarithmic factors) the optimal pseudo-regret guarantee for stochastic rewards and the optimal pseudo-regret or actual regret guarantee for adversarial rewards. Bubeck and Slivkins [BS12] and Auer and Chiang [AC16] begin from a stochastic algorithm and test whether they encounter non-stochastic behavior in which case they switch to adversarial algorithm. In contrast, Seldin et al. [SS14, SL17] begin from an adversarial algorithm with very optimistic learning rate and adapt it if they encounter such behavior. Recently and independently to this work, Wei and Luo [WL18] provide a best of both worlds result with a small-loss pseudo-regret guarantee on the adversarial setting, via a novel analysis of the log-barrier OMD algorithm of Foster et al. [FLL⁺16]. Although the aforementioned algorithms are very elegant, their analysis is not robust to inputs that are slightly away from stochastic. Our work bridges this gap by designing algorithms with a more smooth behavior for close-to-stochastic instances.

There have been other works that attempt to provide improved guarantees than the adversarial setting when instances are well behaved. Hazan and Kale [HK09] offer regret guarantees that scale with the variance of the losses instead of the time horizon. This guarantee is meaningful in settings that have a very predictable nature and have usually the same performance such as routing. However they do not address most applications of stochastic bandits. In Click Fraud, for example the rewards come from Bernoulli

distributions and the variance of such a distribution is high even if the input is totally stochastic. Another approach is the work of Shamir and Szlak [SS17], who consider an input that is adversarial but random local permutations are applied to obtain a more benign instance. This approach is very relevant in settings like buffering, but is again not applicable to our settings.

On the opposite side, attempting to provide improved guarantees for the stochastic setting or enhancing their range is a very active area of research. For instance, the MOSS algorithm [AB09] of Audibert and Bubeck provides the optimal non-distribution-based upper bound for stochastic bandits while retaining the optimal distribution-based stochastic guarantee. The KL-UCB algorithm of Garivier and Cappé [GC11] provides improved constants in the upper bound of the stochastic guarantee matching the lower bound of Lai and Robbins [LR85] for Bernoulli rewards. The Robust UCB algorithm [BCBL13] extends the results to non-bounded rewards replacing with the weaker assumption of bounded variance. However, all the above results are not robust to corruptions from an adaptive adversary due to their deterministic nature. Since the adversary knows the arm the learner will select, they can always corrupt the optimal arm whenever it is about to be selected and therefore cause the learner to either play it multiple times even if it is suboptimal or decide against playing it even with a small amount of corruption (similarly as in our lower bound).

There is also prior work on incorporating corruptions in online decision making. In the online learning front, there are two such attempts, to the best of our knowledge. In their best of both worlds result, Seldin and Slivkins [SS14] allow for some contamination in the data as long as they are obviously selected and they do not decrease the gap by more than a factor of 2. The second work is a recent paper by Gajane et al. [GUK18] who suggest a model of corrupted feedback aiming for differential privacy. Unlike our model, their corruptions are neither adversarial nor adaptive. Both of these works make benign assumptions about the nature of corruption and do not address the main roadblock in the settings we consider: an adversarial saboteur will try to add faulty data in the beginning to change the order between the two arms and, with a minimal corruption, she will achieve this goal. Closer to our model are the works on robust allocation such as online matching with corrupted data [MGZ12, EKM15]; unlike online matching though, in online learning we cannot evaluate the optimum at every round since the algorithm's decisions affect the information it observes.

Last, learning in the presence of corruptions has recently received great attention in the batch learning setting. For instance, recent works study inference under the presence of adversarially corrupted data [MRT15], designing estimators that are robust to corrupted data [DKK⁺16], learning in auctions with some faulty data due to econometrics errors [CD17]. Our work suggests a similar framework for the study of online learning that is robust to adversarial corruptions in the more challenging problem of sequential decision making where decisions also affect the information observed.

2 MODEL

Corrupted stochastic bandits. We study an online bandit learning setting with K arms. Each arm $a \in [K]$ is associated with a distribution $\mathcal{F}(a)$ with mean $\mu(a)$. The distributions are assumed to have positive measure only on rewards in $[0, 1]$ and are unknown to the learner. We refer to the optimal arm as $a^\star = \arg \max_a \mu(a)$ and define $\Delta(a) = \mu(a^\star) - \mu(a)$.¹

We consider an adversary who can corrupt some of the stochastic rewards. The adversary is adaptive, in the sense that the corrupted rewards can be a function of the realization of the stochastic rewards up to that point and of the learner's choices in previous rounds. More formally, the protocol between learner and adversary, at each round $t = 1 \dots T$, is as follows:

- (1) The learner picks a distribution w^t over the K arms.
- (2) Stochastic rewards are drawn for each arm: $r_S^t(a) \sim \mathcal{F}(a)$.
- (3) The adversary observes the realizations of $r_S^t(a)$ as well as rewards and choices of the learner in previous steps and returns a corrupted reward $r^t(a) \in [0, 1]$.
- (4) The learner draws arm $a^t \sim w^t$ and observes $r^t(a^t)$.

We refer to $\max_a |r^t(a) - r_S^t(a)|$ as the amount of corruption injected in round t . The instance is C -corrupted if the total injected corruption is at most C for all realizations of the random variables:

$$\sum_t \max_a |r^t(a) - r_S^t(a)| \leq C$$

Note that the adversary is assumed to be adaptive, in the sense that she has access to all the realizations of random variables for all rounds $\tau < t$ and the realization of rewards at round t but only knows the player's distribution at round t and not the arm a^t . Our guarantees gracefully degrade with the total corruption injected by the adversary.

Regret notions. Regret corresponds to the difference between the reward obtained by the algorithm and the reward of the best arm in hindsight:

$$\text{Reg} = \max_a \sum_t r^t(a) - r^t(a^t)$$

The regret is a random variable that depends on the random rewards, the randomness used by the learner, and the randomness of the adversary. We say that a regret bound $R(T, \delta)$ holds with probability $1 - \delta$ if

$$\mathbb{P}[\text{Reg} < R(T, \delta)] > 1 - \delta$$

where the probability is taken over all the three sources of randomness described.

Finally pseudo-regret is a weaker notion that compares the expected performance of the learner with the arm with the highest expected

¹We note that $a^\star$ is one arm with optimal mean and this does not preclude the existence of other arms with the same mean. If more than one such arms exist, let $a^\star$ be an arbitrary arm with optimal mean and the other arms $a \neq a^\star$ with optimal mean have gap $\Delta(a) = 0$.

performance. In other words:

$$\text{PseudoReg} = \max_a \mathbb{E} \left[\sum_t r^t(a) - r^t(a^t) \right]$$

We note that by Jensen's inequality, $\text{PseudoReg} \leq \mathbb{E}[\text{Reg}]$. We often obtain improved bounds for pseudo-regret since it allows us to offset large positive regret events with large negative regret events.

3 THE UPPER BOUND: MULTI-LAYER ACTIVE ARM ELIMINATION RACE

Active arm elimination. The starting point of our design is the *Active Arm Elimination* algorithm for stochastic bandits [EMM06], which can be viewed as an alternative presentation of the more famous UCB algorithm [ACBF02]. It is based on the following idea: in an initial *exploration phase*, we pull arms in a round-robin fashion and compute an estimate $\tilde{\mu}(a)$ as the average empirical reward of arm a . After $n(a)$ pulls of arm a , usual concentration arguments establish that with probability at least $1 - 1/T^{\Omega(1)}$, the difference of the empirical and actual means is at most $\text{wd}(a) = O(\sqrt{\log(T)/n(a)})$. We say that $[\tilde{\mu}(a) - \text{wd}(a), \tilde{\mu}(a) + \text{wd}(a)]$ is the confidence interval of arm a .

This means in particular that given two arms a and a' , if the difference in empirical means becomes larger than the widths of the confidence intervals, i.e., $\tilde{\mu}(a) - \tilde{\mu}(a') > \text{wd}(a) + \text{wd}(a')$, then with high probability arm a' is not optimal. Once this happens, the algorithm eliminates arm a' by removing it from the round-robin rotation. After both arms a and the optimal arm $a^\star$ are pulled $O(\log(T)/\Delta(a)^2)$ times, the confidence intervals will be small enough that arm a will be eliminated.

Eventually all arms but the optimal are eliminated and we enter what is called the *exploitation phase*. In this phase we only pull the arm with optimal mean. Before we enter exploitation we pulled each suboptimal arm a at most $O(\log(T)/\Delta(a)^2)$ times. Each of those suboptimal pulls incurs regret $\Delta(a)$ in expectation which leads to the pseudo-regret bound of $O(\sum_{a \neq a^\star} \log(T)/\Delta(a))$. This bound can also be converted to a high probability bound if we replace $\log(T)$ by $\log(T/\delta)$.

Arms with small $\Delta(a)$. We note that, for the arms that have $\Delta(a) < 1/\sqrt{T}$, the inverse dependence on the gap may initially seem vacuous; for instance, when there are two optimal arms $a, a^\star$ with the same mean, the upper bound becomes infinite as $\Delta(a) = 0$. However, the inverse dependence on the gap can be replaced by $\Delta(a) \cdot T$ in the case of pseudo-regret and $\sqrt{T}$ in the case of actual regret (due to variance reasons). For simplicity of exposition, we omit this in the current section but we demonstrate how to perform this replacement in Section 5.

3.1 Enlarged Confidence Intervals

The active arm elimination algorithm is clearly not robust to corruption since by corrupting the first $O(\log T)$ steps, the adversary can cause the algorithm to eliminate the optimal arm. As the algorithm never pulls the suboptimal arms after exploration, it is

not able to ever recover. One initial idea to fix this problem is to enlarge the confidence intervals. We can decompose the rewards $r^t(a)$ in two terms $r_S^t(a) + c^t(a)$ where the first term comes from the stochastic reward and the second is the corruption introduced by the adversary. If the total corruption introduced by the adversary is at most C , then with width $\text{wd}(a) = O(\sqrt{\log(T)/n(a)} + C/n(a))$, a similar analysis to above gives us the following regret bound:

Theorem 1. If C is a valid upper bound for the total corruption then active arm elimination with $\text{wd}(a) = \sqrt{\frac{\log(2KT/\delta)}{n(a)}} + \frac{C}{n(a)}$ has regret $O\left(\sum_{a \neq a^*} \left(\frac{\log(KT/\delta) + C}{\Delta(a)}\right)\right)$ with probability $1 - \delta$.

PROOF SKETCH. The proof follows the standard analysis of active arm elimination. We first establish that, with high probability the optimal arm a^* is never inactivated (Lemma 3.1) and then upper bound the number of times each suboptimal arm is played (Lemma 3.2). The pseudo-regret guarantee directly follows by multiplying the number of plays for each arm by its gap $\Delta(a)$. For the high-probability guarantee, we need to also show that the regret incurred in the meantime is not much more than the above. We provide proof details about the theorem and lemmas in the full version of the paper. $\square$

Lemma 3.1. With probability at least $1 - \delta$, arm a^* never becomes inactivated.

Lemma 3.2. With probability at least $1 - \delta$, all arms $a \neq a^*$ become inactivated after $N(a) = \frac{36 \log(2KT/\delta) + 6C}{\Delta(a)^2}$ plays.

3.2 Stochastic Bandits Robust to Known Corruption

The drawback of the active arm elimination algorithm with enlarged confidence intervals (Theorem 1) is that, even if there are no corruptions, it still incurs a regret proportional to C . As a warm up to the main theorem, we provide an algorithm that achieves the usual bound of $O\left(\sum_{a \neq a^*} \frac{\log(KT/\delta)}{\Delta(a)}\right)$ if the input is purely stochastic and, at the same time, achieves $O\left(K \cdot C \cdot \sum_{a \neq a^*} \frac{\log(KT/\delta)^2}{\Delta(a)}\right)$ if the input is C -corrupted for a known C . In the next subsection, we modify the algorithm to make it agnostic to the corruption level C .

Two instances of Active Arm Elimination. The first idea is to run two instances of active arm elimination: the first is supposed to select the correct arm if there is no corruption and the second is supposed to select the right arm if there is C corruption. The first instance is very fast but it is not robust to corruptions. The second instance is slower but more precise, in the sense that it can tolerate corruptions. Since the second instance is more trustworthy, if the second instance decides to eliminate a certain arm a , we eliminate the same arm in the faster instance.

Decrease corruption by sub-sampling. To keep the regret low if the input is stochastic, the second instance of active arm elimination cannot pull a suboptimal arm too many times. Therefore, the technique in Theorem 1 alone is not enough. The *main idea* of the algorithm is to make arm a behave as if it was almost stochastic by

running the second instance with low probability. If the learner selects to run the second instance with probability $1/C$ then, when the adversary adds a certain amount of corruption to a certain round, the second instance observes that corruption with probability $1/C$. Therefore, the expected amount of corruption the learner observes in the second instance is constant. This makes the arms behave almost like stochastic arms in that instance.

Learning algorithm. We obtain our algorithm by combining those ideas. We have two instances of active arm elimination which we denote by F (fast) and S (slow). Each instance keeps an estimate of the mean $\tilde{\mu}^F(a)$ and $\tilde{\mu}^S(a)$ corresponding to the average empirical reward of that arm and also keeps track of how many times each arm was pulled in that instance $n^F(a)$ and $n^S(a)$. This allows us to define a notion of confidence interval in each of the instances. We define $\text{wd}^F(a) = O(\sqrt{\log(T)/n^F(a)})$ as usual and for the slow instance we define slightly larger confidence intervals: $\text{wd}^S(a) = O(\sqrt{\log(T)/n^S(a)} + \log(T)/n^S(a))$ (the reason will be clear in a moment). Also, each instance keeps a set of eliminated arms for that instance: $\mathcal{I}^F$ and $\mathcal{I}^S$.

In each round, with probability $1 - 1/C$ we make a move in the fast instance: we choose the next active arm a in the round robin order, i.e., arm $a \in [K] \setminus \mathcal{I}^F$ which was played less often, pull this arm and increase $n^F(a)$ and update $\tilde{\mu}^F(a)$ accordingly. As usual, if there are two active arms a and a' such that $\tilde{\mu}^F(a) - \tilde{\mu}^F(a') > \text{wd}^F(a) + \text{wd}^F(a')$ we eliminate a' by adding it to $\mathcal{I}^F$.

With the remaining probability we make a move in the slow instance by executing the exact same procedure as described for the other instance. There is only one difference (which causes the two instances to be coupled): when we inactivate an arm a in S we also eliminate it in F. This leaves us with a potential problem: it is possible that all arms in the F instance end up being eliminated. If we reach that point, we play an arbitrary active arm of the slow instance, i.e., any arm $a \in [K] \setminus \mathcal{I}^S$.

The resulting algorithm is formally provided in Algorithm 1.

Algorithm 1 Fast-Slow Active Arm Elimination Race for known corruption C

- 1: Initialize $n^\ell(a) = 0$, $\tilde{\mu}^\ell(a) = 0$, $\mathcal{I}^\ell = \emptyset$ for all $a \in [K]$ and $\ell \in \{F, S\}$
 - 2: **For** Rounds $t = 1..T$
 - 3: Sample algorithm ℓ : $\ell = S$ with probability $1/C$. Else $\ell = F$.
 - 4: **If** $[K] \setminus \mathcal{I}^\ell \neq \emptyset$
 - 5: Play arm $a^t \leftarrow \arg \min_{a \in [K] \setminus \mathcal{I}^\ell} n^\ell(a)$
 - 6: Update $\tilde{\mu}^\ell(a^t) \leftarrow [n^\ell(a) \tilde{\mu}^\ell(a^t) + r^t(a^t)] / [n^\ell(a) + 1]$ and $n^\ell(a) \leftarrow n^\ell(a) + 1$
 - 7: **While** exists arms $a, a' \in [K] \setminus \mathcal{I}^\ell$ with $\tilde{\mu}^\ell(a) - \tilde{\mu}^\ell(a') > \text{wd}^\ell(a) + \text{wd}^\ell(a')$
 - 8: Eliminate a' by adding it to $\mathcal{I}^\ell$
 - 9: **If** $\ell = S$ **then** eliminate a' from the other algorithm by adding it to $\mathcal{I}^F$
 - 10: **Else**
 - 11: Play an arbitrary arm in the set $[K] \setminus \mathcal{I}^S$.
-

Towards the performance guarantee, Lemma 3.3 bounds the amount of corruption that actually enters the slow active arm elimination algorithm, which enables the regret guarantee in Theorem 2.

Lemma 3.3. In Algorithm 1, the slow active arm elimination algorithm S observes, with probability at least $1 - \delta$, corruption of at most $\ln(1/\delta) + 3$ during its exploration phase (when picked with probability $1/C$).

PROOF SKETCH. If one cared just about the expected corruption that affects S, this is at most a constant number since the total corruption is at most C and it affects S with probability $1/C$. To prove a high-probability guarantee against an adaptive adversary, we require a concentration inequality on martingale differences; specifically we apply a Bernstein-style inequality introduced in Lemma 1 of [BLL⁺11]. We provide proof details in the full version of the paper. $\square$

Theorem 2. Algorithm 1 run with widths $\text{wd}^S(a) = \sqrt{\frac{\log(8KT/\delta)}{n^S(a)}} + \frac{2\log(8KT/\delta)}{n^S(a)}$ and $\text{wd}^F(a) = \sqrt{\frac{\log(8KT/\delta)}{n^F(a)}}$ has $O\left(\sum_{a \neq a^*} \frac{\log(KT/\delta)}{\Delta(a)}\right)$ for the stochastic case and $O\left(K \cdot C \cdot \sum_{a \neq a^*} \frac{(\log(KT/\delta))^2}{\Delta(a)}\right)$ for the C -corrupted case with probability at least $1 - \delta$.

PROOF SKETCH. The result for the stochastic case follows standard arguments for stochastic algorithms (since we obtain double the regret of this setting as we run two such algorithms with essentially the same confidence intervals). For the C -corrupted case, we establish via Lemma 3.3 an upper bound on the corruption that will affect the slow active arm elimination algorithm S. Thanks to the sub-sampling, this upper bound is close to a constant instead of depending on C which allows to not incur dependence on C in the stochastic case. Having this upper bound, we can apply it to the algorithm of the previous section to get an upper bound on the number of plays of suboptimal arms in S. Since the algorithms are coupled, such a bound implies an upper bound on the regret that it can cause in F as well. This is because in expectation the arm is played at most $K \cdot C$ times more in F as it may be selected every single time in F prior to getting eliminated by S and F is selected C times more often than S. To obtain the above guarantee with high probability, we lose an extra logarithmic factor. The details of the proof are provided in the full version of the paper. $\square$

3.3 Stochastic Bandits Robust to Agnostic Corruption

Multiple layers of active arm elimination. In the previous subsection we designed an algorithm with two layers: one is faster but cannot tolerate corruptions and the second one is slower but more robust. In order to be agnostic to corruption, we need to plan for all possible amounts of corruption. To achieve this, we introduce $\log(T)$ layers. Each layer is slower but more robust than the previous one. We achieve that by selecting the ℓ -th layer with probability proportional to $2^{-\ell}$. By the argument in the last section, if the corruption level is at most C , then each layer with $\ell \geq \log C$

will observe $O(1)$ corruption in expectation and at most $O(\log T)$ corruption with high probability.

Global eliminations. We couple the $\log T$ instances through what we call global eliminations. If arm a is eliminated by the ℓ -th layer, then we eliminate a in all layers $\ell' \leq \ell$. This is important to prevent us from pulling arm a too often. If arm a is suboptimal and the adversary is C -corrupted, then arm a eventually becomes eliminated in the $\ell^* = \lceil \log C \rceil$ layer after being pulled $\tilde{O}(1/\Delta(a)^2)$ in that layer. Since layer ℓ^* is played with probability $2^{-\ell^*}$ then it takes $\tilde{O}(C/\Delta(a)^2)$ iterations until arm is eliminated globally, in which case we will have total regret at most $\tilde{O}(C/\Delta(a))$ from that arm.

Multi-layer active arm elimination race. We now describe our main algorithm in the paper. We call it a *race* since we view it as multiple layers racing to pick the optimal arm. The less robust layers are faster so they arrive first and we keep choosing (mostly) according to them until more robust but slower layers finish and correct or confirm the current selection of the best arm.

The algorithm keeps $\ell = 1 \dots \log(T)$ different instances of active arm elimination. The ℓ -th instance has as state the empirical means of each arm $\tilde{\mu}^\ell(a)$, the number $n^\ell(a)$ of times each arm a was pulled and the set I^ℓ of inactive arms. The width of the confidence interval for arm a in the ℓ -th layer is implicitly defined as $\text{wd}^\ell(a) = O(\sqrt{\log(T)/n^\ell(a)} + \log(T)/n^\ell(a))$.

In each round t we sample $\ell \in \{1, \dots, \log T\}$ with probability $2^{-\ell}$ (with remaining probability we pick layer 1). When layer ℓ is selected, we make a move in the active arm elimination instance corresponding to that layer: we sample the active arm in that layer with the least number of pulls, i.e., arm $a \in [K] \setminus I^\ell$ minimizing $n^\ell(a)$. In case $[K] \setminus I^\ell$ is empty, we pull an arbitrary arm from $[K] \setminus I^{\ell'}$ for the lowest ℓ' such that $[K] \setminus I^{\ell'}$ is non-empty.

The way we couple different layers is that once arm a' is eliminated in layer ℓ because there is another active arm a in layer ℓ such that $\tilde{\mu}^\ell(a) - \tilde{\mu}^\ell(a') < \text{wd}^\ell(a) + \text{wd}^\ell(a')$ we eliminate arm a' in all previous layers, keeping the invariant that: $I^1 \supseteq I^2 \supseteq I^3 \supseteq \dots$

Figure 1 provides an example of the state of the algorithm, which is formally defined in Algorithm 2.

We now provide the main result of the paper, a regret guarantee for Algorithm 2:

Theorem 3. Algorithm 2 which is agnostic to the corruption level C , when run with widths $\text{wd}^\ell(a) = \sqrt{\frac{\log(4KT \cdot \log T/\delta)}{n^\ell(a)}} + \frac{\log(4KT \cdot \log T/\delta)}{n^\ell(a)}$ has regret:

$$O\left(\sum_{a \neq a^*} \frac{K \cdot C \log(KT/\delta) + \log(T)}{\Delta(a)} \cdot \log(KT/\delta)\right).$$

PROOF SKETCH. Similarly to the previous theorem, the regret guarantee comes from the summation between layers that are essentially stochastic (where the corruption is below their corruption level, i.e. less than $C \leq 2^\ell$ for layer ℓ). From each of these layers, we incur $O\left(\frac{\log(KT/\delta)}{\Delta(a)}\right)$ regret. Since there are at most $\log(T)$ such layers, the second term in the theorem is derived. The challenge is to bound

	arm 1	arm 2	...	arm d
$\ell = 1$	$\tilde{\mu}^1(1), n^1(1)$	$\tilde{\mu}^1(2), n^1(2)$	...	$\tilde{\mu}^1(d), n^1(d)$
$\ell = 2$	$\tilde{\mu}^2(1), n^2(1)$	$\tilde{\mu}^2(2), n^2(2)$	...	$\tilde{\mu}^2(d), n^2(d)$
$\ell = 3$	$\tilde{\mu}^3(1), n^3(1)$	$\tilde{\mu}^3(2), n^3(2)$	...	$\tilde{\mu}^3(d), n^3(d)$
$\vdots$	$\vdots$	$\vdots$		$\vdots$
$\ell = \lg T$	$\tilde{\mu}^{\lg T}(1), n^{\lg T}(1)$	$\tilde{\mu}^{\lg T}(2), n^{\lg T}(2)$	...	$\tilde{\mu}^{\lg T}(d), n^{\lg T}(d)$

Figure 1: Example of the state of the algorithm: for each layer ℓ and arm a we keep the estimated mean $\tilde{\mu}^\ell(a)$ and the number of pulls $n^\ell(a)$. Red cells indicate arms that have been eliminated in that layer. If an arm is eliminated in a layer, it is eliminated in all previous layers. If a layer where all the arms are eliminated (like layer 1 in the figure) is selected, we play an arbitrary active arm with the lowest layer that contains active arms.

Algorithm 2 Multi-layer Active Arm Elimination Race

- 1: Initialize $n^\ell(a) = 0$, $\tilde{\mu}^\ell(a) = 0$, $\mathcal{I}^\ell = \emptyset$ for all $a \in [K]$ and $\ell \in [\log T]$
 - 2: **For** Rounds $t = 1..T$
 - 3: Sample layer $\ell \in [\log T]$ with probability $2^{-\ell}$. With remaining prob, sample $\ell' = 1$
 - 4: **If** $[K] \setminus \mathcal{I}^{\ell'} \neq \emptyset$
 - 5: Play arm $a^t \leftarrow \arg \min_{a \in [K] \setminus \mathcal{I}^{\ell'}} n^{\ell'}(a)$
 - 6: Update $\tilde{\mu}^{\ell'}(a^t) \leftarrow [n^{\ell'}(a^t)\tilde{\mu}^{\ell'}(a^t) + r^t(a^t)]/[n^{\ell'}(a^t) + 1]$ and $n^{\ell'}(a^t) \leftarrow n^{\ell'}(a^t) + 1$
 - 7: **While** exists arms $a, a' \in [K] \setminus \mathcal{I}^{\ell'}$ with $\tilde{\mu}^{\ell'}(a) - \tilde{\mu}^{\ell'}(a') > \text{wd}^{\ell'}(a) + \text{wd}^{\ell'}(a')$
 - 8: Eliminate a' by adding it to $\mathcal{I}^{\ell'}$ for all $\ell' \leq \ell$
 - 9: **Else**
 - 10: Find minimum ℓ' such that $[K] \setminus \mathcal{I}^{\ell'} \neq \emptyset$ and play an arbitrary arm in that set.
-

the regret incurred by layers that are not robust to the corruption. However, there exists some layer ℓ^* that is above the corruption level. By bounding the amount of steps that this level will require in order to inactivate each arm $a \neq a^*$ in the incorrect layers (via Lemma 3.2), we obtain similarly to Theorem 2 a bound on the regret caused by this arm in those layers. Since we take the minimum such layer and the tolerance of layers is within powers of 2, the fact that its corruption level does not match exactly the corruption that occurred only costs an extra factor of 2 in the regret. The details of the proof are provided in the full version of the paper. $\square$

4 THE LOWER BOUND

For the two arms case where the gap between the arms is $\Delta > 0$, Theorem 2 presents an algorithm which achieves $\mathcal{O}(\log T/\Delta)$ pseudo-regret if the input is stochastic and $\mathcal{O}(C \log(T/\delta)/\Delta)$ with probability $1 - \delta$ if the input is at most C -corrupted. We show below that this dependence is tight.

The lower bound (Theorem 4) adapts the technique of Auer and Chiang [AC16] from the adversarial to the corrupted setting. The main idea is that an algorithm with logarithmic regret in the stochastic setting cannot query the sub-optimal arm more than $\log(T)/\Delta^2$ times. This implies a long time period where the learner queries the input only constant number of times. By corrupting all rounds before this period, an adversary can make the optimal arm look sub-optimal and trick the learner into not pulling the optimal arm for long time, causing large regret. Theorem 5 adapts this argument bounding the expected positive regret $\mathbb{E}[\text{Reg}^+]$ where $x^+ = \max\{x, 0\}$; the high probability bounds provided imply bounds on the expected positive regret. Both proofs are provided in the full version of the paper.

Theorem 4. Consider a multi-armed bandits algorithm that has the property that for any stochastic input in the two arm setting, it has pseudo-regret bounded by $c \log(T)/\Delta$, where $\Delta = |\mu_1 - \mu_2|$. For any $\epsilon, \epsilon' \in (0, 1)$, there is a corruption level C with $T^\epsilon < C < T^{\epsilon'}$ and a C -corrupted instance such that with constant probability the regret is $\Omega(C)$.

Theorem 5. If a multi-armed bandits algorithm that has the property that for any stochastic input in the two arm setting, it has pseudo-regret bounded by $c \log^{1+\alpha}(T)/\Delta$ for $\alpha < 1$. For any $\epsilon, \epsilon' \in (0, 1)$, there is a corruption level C with $T^\epsilon < C < T^{\epsilon'}$ and a C -corrupted instance such that $\mathbb{E}[\text{Reg}^+] = \Omega(T^{\epsilon-\delta})$ for all $\delta > 0$.

5 EXTENSIONS

In this section, we discuss some extensions that our algorithm can accommodate.

Definition of corruption. We presented all results measuring the corruption as the sum over all rounds of the maximum across arms of the corruption injected by the adversary:

$$\sum_t \max_a |r^t(a) - r_S^t(a)| \leq C.$$

In fact all our results can be improved via using $C(a) = \sum_t |r^t(a) - r_S^t(a)|$ and replacing C by $\max(C(a), C(a^*))$ for summand a . More formally, our main theorem (Theorem 3) becomes:

Theorem 6. Algorithm 2 which is agnostic to the corruptions $C(a) = \sum_t |r^t(a) - r_S^t(a)|$ has regret of

$$\mathcal{O}\left(\sum_{a \neq a^*} \frac{K \cdot \max(C(a^*), C(a)) \cdot \log(KT/\delta) + \log(T)}{\Delta(a)} \cdot \log(KT/\delta)\right)$$

when run with widths $\text{wd}^\ell(a) = \sqrt{\frac{\log(4KT \cdot \log T/\delta)}{n^\ell(a)}} + \frac{\log(4KT \cdot \log T/\delta)}{n^\ell(a)}$.

The proof follows the same arguments since it only compares each arm a with a^* . This result is nice since the contribution of each arm to the regret is a function only of its own gap and the corruption injected to it and the one injected to arm a^* . The latter dependence on the corruption on the optimal arm is essential since the main attack we presented to the classical arguments only corrupts arm a^* – the lower bound of the previous section also only adds corruption to a^* .

Dependence on the gap. In Section 3, all our guarantees have an inverse dependence on the gap $\Delta(a)$ of all arms a . Note that such a guarantee is completely meaningless for arms with a very small gap; for instance, if there exist two optimal arms then there is an arm $a \neq a^*$ with $\Delta(a) = 0$ which makes the presented bound infinite and therefore vacuous. As we hinted there though, this inverse dependence can be improved for arms with small $\Delta(a) \leq 1/\sqrt{T}$. Our proofs generally relied on setting an upper bound on the number of times that a suboptimal arm is played and thereby providing an upper bound on the regret they cause.

For arms with $\Delta(a) \leq 1/\sqrt{T}$, an alternative analysis is to say that, even if they are erroneously selected every single time, we can upper bound the loss in performance they cause. For pseudo-regret, the performance loss if they were selected every single time is $\Delta(a) \cdot T \leq \sqrt{T} \leq \log T / \Delta(a)$. For actual regret, one needs to also take into consideration the variance but, even if they are selected every single time, a Hoeffding bound shows that their total reward is with high probability at most $\sqrt{T}$ lower than its expectation. As a result, the inverse dependence on $\Delta(a)$ in our bound can be replaced by $\min(\Delta(a) \cdot T, 1/\Delta(a))$ for pseudo-regret and $\min(\sqrt{T}, 1/\Delta(a))$ for actual regret.

Moreover, the careful reader may have noticed that in Theorem 1, the dependence $\frac{C}{\Delta(a)}$ can be replaced by a sole dependence on C without the gap. However, this does not extend to the subsequent theorems since the dependence on C there does not come from the upper bound on the corruption experienced (this is at most $\log T$ due to subsampling). Instead, the dependence on C comes from projecting the correct layer (smallest layer robust to corruption) to the previous layers via the number of times it will take to eliminate any suboptimal arm.

Uncorrupted objective. In applications such as spam, the corruptions should not be counted as part of the rewards. Our algorithm provides the same guarantee in the case of uncorrupted rewards (the difference between the performances in the two objectives is at most C which is already in the bound). One can also observe that the linear dependence on C is still necessary: consider 2 arms with $\Delta = 1$ and an adversary that corrupts the first C steps making them look identical. The learner has no better option than randomly selecting between the two which gives him a regret of $C/2$ under the uncorrupted objective. We note that, in this setting, the linear dependence is necessary unconditionally of the performance of the algorithm in the stochastic setting.

Towards best of all worlds. In the previous section, we showed that a logarithmic dependence in the stochastic setting comes at the expense of linear dependence on C in the C -corrupted setting if we focus on actual regret. A very interesting direction is to achieve such an improvement with either a higher power on the logarithm in the stochastic setting or aiming for pseudo-regret instead.

In fact, we can combine our algorithm with the SAPO algorithm of Auer and Chiang [AC16] and achieve a bicriteria guarantee for pseudo-regret. For an $a < 1/2$ specified by the algorithm, we achieve our guarantee if the corruption is $C \leq T^a$ and at most $T^{1/2+a}$ otherwise; notice that the case $a = 0$ corresponds to the best of both worlds. This is done via running the SAPO algorithm at the

level $a \log(T)$ with probability T^{-a} instead of having higher layers. The SAPO algorithm guarantees that the pseudo-regret caused by any particular arm is at most logarithmic if the instance is stochastic and at most $\sqrt{T}$ if it is adversarial via a beautiful analysis that keeps negative regret of time intervals that have performed well to avoid testing eliminated arms too often. In our setting, if the corruption level is less than T^a , the instance behaves as stochastic causing at most logarithmic regret. Else the instance is corrupted and we can extrapolate the regret in this layer to the whole algorithm as arms that are eliminated in this layer are also eliminated before via global eliminations. Since the regret there is at most $\sqrt{T}$ and this is multiplied by T^a , this implies a bound of $T^{1/2+a}$ on pseudo-regret.

FULL VERSION

The full version of the paper can be found in the following link: <https://arxiv.org/abs/1803.09353>. The reader is advised to consult this for proof details and updated versions of the paper.

ACKNOWLEDGEMENTS

The first author was supported under NSF grant CCF-1563714. Part of the work was done while the author was interning at Google. The authors would like to thank Sid Banerjee whose lecture notes on stochastic bandits proved very helpful, Andrés Munoz Medina, Karthik Sridharan, and Éva Tardos for useful discussions, Manish Raghavan for suggestions on the writeup, and the anonymous reviewers for the valuable feedback they provided that improved the presentation of the paper.

REFERENCES

- [AB09] Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*, 2009.
- [AC16] Peter Auer and Chao-Kai Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*, 2016.
- [ACBF02] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.*, 47(2-3):235–256, May 2002.
- [ACBFS03] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77, January 2003.
- [BCB12] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [BCBL13] Sébastien Bubeck, Nicolò Cesa-Bianchi, and Gábor Lugosi. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013.
- [BLL⁺11] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert E. Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [BS12] Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*, 2012.
- [CBL06] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.
- [CD17] Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2017.
- [DKK⁺16] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *Proceedings of the 57th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2016.

- [EKM15] Hossein Esfandiari, Nitish Korula, and Vahab S. Mirrokni. Online allocation with traffic spikes: Mixing adversarial and stochastic models. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, 2015.
- [EMM06] 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*, 7:1079–1105, 2006.
- [FLL⁺16] Dylan J Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [GC11] Aurélien Garivier and Olivier Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *24th Annual Conference on Learning Theory (COLT)*, 2011.
- [GUK18] Pratik Gajane, Tanguy Urvoy, and Emilie Kaufmann. Corrupt bandits for privacy preserving input. In *29th International Conference on Algorithmic Learning Theory (ALT)*, 2018.
- [HK09] Elad Hazan and Satyen Kale. Better algorithms for benign bandits. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2009.
- [LR85] T.L. Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.*, 6(1):4–22, March 1985.
- [MGZ12] Vahab S. Mirrokni, Shayan Oveis Gharan, and Morteza Zadimoghaddam. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2012.
- [MRT15] Yishay Mansour, Aviad Rubinstein, and Moshe Tennenholtz. Robust probabilistic inference. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015.
- [SL17] Yevgeny Seldin and Gábor Lugosi. An improved parametrization and analysis of the EXP3++ algorithm for stochastic and adversarial bandits. In *Proceedings of the 30th Conference on Learning Theory (COLT)*, 2017.
- [Sli17] Aleksandrs Slivkins. Introduction to multi-armed bandits. 2017.
- [SS14] Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *Proceedings of the 31th International Conference on Machine Learning (ICML)*, 2014.
- [SS17] Ohad Shamir and Liran Szlak. Online learning with local permutations and delayed feedback. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017.
- [WL18] Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. *CoRR*, abs/1801.03265, 2018.