
One Practical Algorithm for Both Stochastic and Adversarial Bandits

Yevgeny Seldin

Queensland University of Technology, Brisbane, Australia

YEVGENY.SELDIN@GMAIL.COM

Aleksandrs Slivkins

Microsoft Research, New York NY, USA

SLIVKINS@MICROSOFT.COM

Abstract

We present an algorithm for multiarmed bandits that achieves almost optimal performance in both stochastic and adversarial regimes without prior knowledge about the nature of the environment. Our algorithm is based on augmentation of the EXP3 algorithm with a new control lever in the form of exploration parameters that are tailored individually for each arm. The algorithm simultaneously applies the “old” control lever, the learning rate, to control the regret in the adversarial regime and the new control lever to detect and exploit gaps between the arm losses. This secures problem-dependent “logarithmic” regret when gaps are present without compromising on the worst-case performance guarantee in the adversarial regime. We show that the algorithm can exploit both the usual expected gaps between the arm losses in the stochastic regime and deterministic gaps between the arm losses in the adversarial regime. The algorithm retains “logarithmic” regret guarantee in the stochastic regime even when some observations are contaminated by an adversary, as long as on average the contamination does not reduce the gap by more than a half. Our results for the stochastic regime are supported by experimental validation.

1. Introduction

Stochastic multiarmed bandits (Thompson, 1933; Robbins, 1952; Lai & Robbins, 1985; Auer et al., 2002a) and adversarial multiarmed bandits (Auer et al., 1995; 2002b) have co-existed in parallel for almost two decades by now, in the sense that no algorithm for stochastic multiarmed bandits is applicable to adversarial multiarmed bandits and al-

gorithms for adversarial bandits are unable to exploit the simpler regime of stochastic bandits. The recent attempt of Bubeck & Slivkins (2012) to bring them together did not make it in the full sense of unification, since the algorithm of Bubeck and Slivkins relies on the knowledge of time horizon and makes a one-time irreversible switch between stochastic and adversarial operation modes if the beginning of the game is estimated to exhibit adversarial behavior.

We present an algorithm that treats both stochastic and adversarial multiarmed bandit problems without distinguishing between them. Our algorithm “just runs”, as most other bandit algorithms, without knowledge of time horizon and without making any hard statements about the nature of the environment. We show that if the environment happens to be adversarial the performance of the algorithm is just a factor of 2 worse than the performance of the EXP3 algorithm (with the best constants, as described in Bubeck & Cesa-Bianchi (2012)) and if the environment happens to be stochastic the performance of our algorithm is comparable to the performance of UCB1 of Auer et al. (2002a). Thus, we cover the full range and achieve almost optimal performance at the extreme points.

Furthermore, we show that the new algorithm can exploit both the usual expected gaps between the arm losses in the stochastic regime and deterministic gaps between the arm losses in the adversarial regime. We also show that the algorithm retains “logarithmic” regret guarantee in the stochastic regime even when some observations are adversarially contaminated, as long as on average the contamination does not reduce the gap by more than a half. To the best of our knowledge, no other algorithm has been yet shown to be able to exploit gaps in the adversarial or adversarially contaminated stochastic regimes. The contaminated stochastic regime is a very practical model, since in many real-life situations we are dealing with stochastic environments with occasional disturbances.

Since the introduction of Thompson’s sampling (Thompson, 1933) (which was analyzed only after 80 years (Kaufmann et al., 2012; Agrawal & Goyal, 2013)) a variety of al-

gorithms were invented for the stochastic multiarmed bandit problem. The most powerful for today are KL-UCB (Cappé et al., 2013), EwS (Maillard, 2011), and the aforementioned Thompson’s sampling. It is easy to show that any deterministic algorithm can potentially suffer linear regret in the adversarial regime (see the supplementary material for a proof). Although nothing is known about the performance of randomized algorithms for stochastic bandits in the adversarial regimes, empirically they are extremely sensitive to deviations from the stochastic assumption.

In the adversarial world the most powerful algorithm for today is INF (Audibert & Bubeck, 2009; Bubeck & Cesa-Bianchi, 2012). Nevertheless, the EXP3 algorithm of Auer et al. (2002b) still retains an important place, mainly due to its simplicity and wide applicability, which covers combinatorial bandits, partial monitoring games, and many other adversarial problems. Since any stochastic problem can be seen as an instance of an adversarial problem, both INF and EXP3 have the worst-case “root- t ” regret guarantee in the stochastic regime, but it is not known whether they can do better. Empirically in the stochastic regime EXP3 is inferior to all other known algorithms for this setting, including the simplest UCB1 algorithm.

It is interesting to take a brief look into the development of EXP3. The algorithm was first suggested in Auer et al. (1995) and its parametrization and analysis were improved in Auer et al. (2002b). The EXP3 of Auer et. al. was designed for the multiarmed bandit game with rewards and its playing strategy is based on mixing Gibbs distribution (also known as “exponential weights”) with a uniform exploration distribution in proportion to the learning rate. The uniform exploration leaves no hope for achieving “logarithmic” regret in the stochastic regime simultaneously with the “root- t ” regret in the adversarial regime, since each arm is played at least $\Omega(\sqrt{t})$ times in t rounds of the game. By changing the learning rate Cesa-Bianchi & Fischer (1998) managed to derive a different parametrization of the algorithm that was shown to achieve “logarithmic” regret in the stochastic regime, but it had no regret guarantees in the adversarial regime. Stoltz (2005) has observed that in the game with losses the “root- t ” regret guarantee in the adversarial regime can be achieved without mixing in the uniform distribution (and even lead to better constants).¹ However, mixing in any distribution that element-wise does not exceed the learning rate does not break the worst-case performance of the algorithm in the game with losses. We exploit this emerged freedom in order to derive a modification of the EXP3 algorithm that achieves almost optimal regret in both adversarial and stochastic regimes without prior knowledge about the nature of the environment.

¹Rewards can be transformed into losses by taking $\ell = 1 - r$.

2. Problem Setting

We study the multiarmed bandit (MAB) game with losses. In each round t of the game the algorithm chooses one action A_t among K possible actions, a.k.a. *arms*, and observes the corresponding loss $\ell_t^{A_t}$. The losses of other arms are not observed. There is a large number of loss generation models, four of which are considered below. In this work we restrict ourselves to loss sequences $\{\ell_t^a\}_{t,a}$ that are generated independently of the algorithm’s actions. Under this assumption we can assume that the loss sequences are written down before the game starts (but not revealed to the algorithm). We also make a standard assumption that the losses are bounded in the $[0, 1]$ interval.

The performance of the algorithm is quantified by *regret*, defined as the difference between the expected loss of the algorithm up to round t and the expected loss of the best arm up to round t :

$$R(t) = \sum_{s=1}^t \mathbb{E} [\ell_s^{A_s}] - \min_a \left\{ \mathbb{E} \left[\sum_{s=1}^t \ell_s^a \right] \right\}.$$

The expectation is taken over the possible randomness of the algorithm and loss generation model. The goal of the algorithm is to minimize the regret.

We consider two standard loss generation models, the *adversarial regime* and the *stochastic regime* and two intermediate regimes, the *contaminated stochastic regime* and the *adversarial regime with a gap*.

Adversarial regime. In this regime the loss sequences are generated by an unrestricted adversary (who is oblivious to the algorithm’s actions). This is the most general setting and the other three regimes can be seen as special cases of the adversarial regime. An arm $a \in \arg \min_{a'} \left(\sum_{s=1}^t \ell_s^{a'} \right)$ is known as a *best arm in hindsight* for the first t rounds.

Stochastic regime. In this regime the losses ℓ_t^a are sampled independently from an unknown distribution that depends on a , but not on t . We use $\mu(a) = \mathbb{E} [\ell_t^a]$ to denote the expected loss of arm a . Arm a is called a *best arm* if $\mu(a) = \min_{a'} \{\mu(a')\}$ and *suboptimal* otherwise; let a^* denote some best arm. For each arm a , define the *gap* $\Delta(a) = \mu(a) - \mu(a^*)$. Let $\Delta = \min_{a: \Delta(a) > 0} \{\Delta(a)\}$ denote the minimal gap.

Letting $N_t(a)$ be the number of times arm a was played up to (and including) round t , the regret can be rewritten as

$$R(t) = \sum_a \mathbb{E} [N_t(a)] \Delta(a). \quad (1)$$

Contaminated stochastic regime. In this regime the adversary picks some round-arm pairs (t, a) (“locations”) before the game starts and assigns the loss values there in an

arbitrary way. The remaining losses are generated according to the stochastic regime.

We call a contaminated stochastic regime *moderately contaminated after τ rounds* if for all $t \geq \tau$ the total number of contaminated locations of each suboptimal arm up to time t is at most $t\Delta(a)/4$ and the number of contaminated locations of each best arm is at most $t\Delta/4$. By this definition, for all $t \geq \tau$ on average (over stochasticity of the loss sequences) the adversary can reduce the gap of every arm by at most a half.

Adversarial regime with a gap. An adversarial regime is named by us *an adversarial regime with a gap* if there exists a round τ and an arm a_τ^* that persists to be the best arm in hindsight for all rounds $t \geq \tau$. We name such arm a *consistently best arm after round τ* . If no such arm exists then a_τ^* is undefined. Note that if a_τ^* is defined for some τ then $a_{\tau'}^*$ is defined for all $\tau' > \tau$.

We use $\lambda_t(a) = \sum_{s=1}^t \ell_s^a$ to denote the cumulative loss of arm a . Whenever a_τ^* is defined we define a *deterministic gap* of arm a on round τ as:

$$\Delta(\tau, a) = \min_{t \geq \tau} \left\{ \frac{1}{t} (\lambda_t(a) - \lambda_t(a_\tau^*)) \right\}.$$

If a_τ^* is undefined, $\Delta(\tau, a)$ is defined as zero.

Notation. We use $\mathbb{1}_{\{E\}}$ to denote the indicator function of event E and $\mathbb{1}_t^a = \mathbb{1}_{\{A_t=a\}}$ to denote the indicator function of the event that arm a was played on round t .

3. Main Results

Our main results include a new algorithm, which we name EXP3++, and its analysis in the four regimes defined in the previous section. The EXP3++ algorithm, provided in Algorithm 1 box, is a generalization of the EXP3 algorithm with losses.

Algorithm 1 Algorithm EXP3++.

Remark: See text for definition of η_t and $\xi_t(a)$.

$\forall a: \tilde{L}_0(a) = 0.$

for $t = 1, 2, \dots$ **do**

$$\beta_t = \frac{1}{2} \sqrt{\frac{\ln K}{tK}}.$$

$$\forall a: \varepsilon_t(a) = \min \left\{ \frac{1}{2K}, \beta_t, \xi_t(a) \right\}.$$

$$\forall a: \rho_t(a) = e^{-\eta_t \tilde{L}_{t-1}(a)} / \sum_{a'} e^{-\eta_t \tilde{L}_{t-1}(a')}.$$

$$\forall a: \tilde{\rho}_t(a) = (1 - \sum_{a'} \varepsilon_t(a')) \rho_t(a) + \varepsilon_t(a).$$

Draw action A_t according to $\tilde{\rho}_t$ and play it.

Observe and suffer the loss $\ell_t^{A_t}$.

$$\forall a: \tilde{\ell}_t^a = \frac{\ell_t^{A_t}}{\tilde{\rho}_t(a)} \mathbb{1}_t^a.$$

$$\forall a: \tilde{L}_t(a) = \tilde{L}_{t-1}(a) + \tilde{\ell}_t^a.$$

end for

The EXP3++ algorithm has two control levers: the *learn-*

ing rate η_t and the *exploration parameters* $\xi_t(a)$. The EXP3 with losses (as described in Bubeck & Cesa-Bianchi (2012)) is a special case of the EXP3++ with $\eta_t = 2\beta_t$ and $\xi_t(a) = 0$.

The crucial innovation in EXP3++ is the introduction of exploration parameters $\xi_t(a)$, which are tuned individually for each arm depending on the past observations. In the sequel we show that tuning only the learning rate η_t suffices to control the regret of EXP3++ in the adversarial regime, irrespective of the choice of the exploration parameters $\xi_t(a)$. Then we show that tuning only the exploration parameters $\xi_t(a)$ suffices to control the regret of EXP3++ in the stochastic regime irrespective of the choice of η_t , as long as $\eta_t \geq \beta_t$. Applying the two control levers simultaneously we obtain an algorithm that achieves the optimal “root- t ” regret in the adversarial regime (up to logarithmic factors) and almost optimal “logarithmic” regret in the stochastic regime (though with a suboptimal power in the logarithm). Then show that the new control lever is even more powerful and allows to detect and exploit the gap in even more challenging situations, including moderately contaminated stochastic regime and adversarial regime with a gap.

Adversarial Regime

First, we show tuning η_t is sufficient to control the regret of EXP3++ in the adversarial regime.

Theorem 1. *For $\eta_t = \beta_t$ and any $\xi_t(a) \geq 0$ the regret of EXP3++ for any t satisfies:*

$$R(t) \leq 4\sqrt{Kt \ln K}.$$

Note that the regret bound in Theorem 1 is just a factor of 2 worse than the regret of EXP3 with losses (Bubeck & Cesa-Bianchi, 2012).

Stochastic Regime

Now we show that for any $\eta_t \geq \beta_t$ tuning the exploration parameters $\xi_t(a)$ suffices to control the regret of the algorithm in the stochastic regime. By choosing $\eta_t = \beta_t$ we obtain algorithms that have both the optimal “root- t ” regret scaling in the adversarial regime *and* “logarithmic” regret scaling in the stochastic regime.

We consider a number of different ways of tuning the exploration parameters $\xi_t(a)$, which lead to different parametrizations of EXP3++. We start with an idealistic assumption that the gap is known, just to give an idea of what is the best result we can hope for.

Theorem 2. *Assume that the gaps $\Delta(a)$ are known. For any choice of $\eta_t \geq \beta_t$ and any $c \geq 18$, the regret of EXP3++ with $\xi_t(a) = \frac{c \ln(t\Delta(a)^2)}{t\Delta(a)^2}$ in the stochastic regime*

satisfies:

$$R(t) \leq \sum_a O\left(\frac{\ln(t)^2}{\Delta(a)}\right) + \sum_a \tilde{O}\left(\frac{K}{\Delta(a)^3}\right).$$

The constants in this theorem are small and are provided explicitly in the analysis. We also show that c can be made almost as small as 2.

Next we show that using the empirical gap as an estimate of the true gap

$$\hat{\Delta}_t(a) = \min\left\{1, \frac{1}{t} \left(\tilde{L}_t(a) - \min_{a'} \left(\tilde{L}_t(a')\right)\right)\right\} \quad (2)$$

we can also achieve polylogarithmic regret guarantee. We call this algorithm EXP3++^{AVG}.

Theorem 3. *Let $c \geq 18$ and $\eta_t \geq \beta_t$. Let t^* be the minimal integer that satisfies $t^* \geq \frac{4c^2 K \ln(t^*)^4}{\ln(K)}$ and let $t^*(a) = \max\left\{t^*, \left\lceil e^{1/\Delta(a)^2} \right\rceil\right\}$. The regret of EXP3++ with $\xi_t(a) = \frac{c(\ln t)^2}{t\hat{\Delta}_{t-1}(a)^2}$ (termed EXP3++^{AVG}) in the stochastic regime satisfies:*

$$R(t) \leq \sum_a O\left(\frac{\ln(t)^3}{\Delta(a)}\right) + \sum_a \Delta(a)t^*(a).$$

Although the additive constants $t^*(a)$ in this theorem are very large, in the experimental section we show that a minor modification of this algorithm performs comparably to UCB1 in the stochastic regime (and has the adversarial regret guarantee in addition).

In the following theorem we show that if we assume a known time horizon T , then we can eliminate the additive term $e^{1/\Delta(a)^2}$ in the regret bound. The algorithm in Theorem 4 replaces the empirical gap estimate (2) in the definition of $\xi_t(a)$ with a lower confidence bound on the gap and slightly adjusts other terms. We name this algorithm EXP3++^{LCBT}.

Theorem 4. *Consider the stochastic regime with a known time horizon T . The EXP3++^{LCBT} algorithm with any $\eta_t \geq \beta_t$ and appropriately defined $\xi_t(a)$ achieves regret $R(T) \leq \frac{O(\log^3 T)}{\Delta^3}$.*

The precise definition of EXP3++^{LCBT} and the proof of Theorem 4 are provided in the supplementary material. It seems that simultaneous elimination of the assumption on the known time horizon and the exponentially large additive term is a very challenging problem and we defer it for future work.

Contaminated Stochastic Regime

Next we show that EXP3++^{AVG} can sustain moderate contamination in the stochastic regime without a significant deterioration in performance.

Theorem 5. *Under the parametrization given in Theorem 3, for $t^*(a) = \max\left\{t^*, \left\lceil e^{4/\Delta(a)^2} \right\rceil\right\}$, where t^* is defined as before, the regret of EXP3++^{AVG} in the stochastic regime that is moderately contaminated after τ rounds satisfies:*

$$R(t) \leq \sum_a O\left(\frac{\ln(t)^3}{\Delta(a)}\right) + \sum_a \max\{t^*(a), \tau\}.$$

The price that is paid for moderate contamination after τ rounds is the scaling of $\Delta(a)$ by a factor of 1/2 and the additive factor of τ . (The scaling of Δ affects the definition of t^* and the constant in $O\left(\frac{\ln(t)^3}{\Delta(a)}\right)$.) As before, the regret guarantee of Theorem 5 comes in addition to the guarantee of Theorem 1.

Adversarial Regime with a Gap

Finally, we show that EXP3++^{AVG} can also take advantage of deterministic gap in the adversarial regime.

Theorem 6. *Under the parametrization given in Theorem 3, the regret of EXP3++^{AVG} in the adversarial regime satisfies:*

$$R(t) \leq \sum_a \min_{\tau} \left\{ \max\left\{t^*, \tau, e^{1/(\Delta(\tau, a))^2}\right\} + O\left(\frac{\ln(t)^3}{\Delta(\tau, a)}\right) \right\}.$$

We remind the reader that in the absence of consistently best arm $\Delta(\tau, a)$ is defined as zero and the regret bound is vacuous (but the regret bound of Theorem 1 still holds). We also note that $\Delta(\tau, a)$ is a non-decreasing function of τ . Therefore, there is a trade-off: increasing τ increases $\Delta(\tau, a)$, but loses the regret guarantee on the rounds before τ (for simplicity, we assume that we have no guarantees before τ). Theorem 6 allows to pick τ that minimizes this trade-off. An important implication of the theorem is that if the deterministic gap is growing with time the regret guarantee improves too.

4. Proofs

We prove the theorems from the previous section in the order they were presented.

The Adversarial Regime

The proof of Theorem 1 relies on the following lemma, which is an intermediate step in the analysis of EXP3 by Bubeck (2010) (see also Bubeck & Cesa-Bianchi (2012)).

Lemma 7. *For any K sequences of non-negative numbers X_1^a, X_2^a, \dots indexed by $a \in \{1, \dots, K\}$ and any non-increasing positive sequence η_1, η_2, \dots , for $\rho_t(a) =$*

$\frac{\exp(-\eta_t \sum_{s=1}^{t-1} X_s^a)}{\sum_{h'} \exp(-\eta_t \sum_{s=1}^{t-1} X_s^a)}$ (assuming for $t = 1$ the sum in the exponent is zero) we have:

$$\begin{aligned} & \sum_{t=1}^T \sum_a \rho_t(a) X_t^a - \min_a \left(\sum_{t=1}^T X_t^a \right) \\ & \leq \frac{1}{2} \sum_{t=1}^T \eta_t \sum_a \rho_t(a) (X_t^a)^2 + \frac{\ln K}{\eta_T}. \end{aligned} \quad (3)$$

More precisely, we are using the following corollary, which follows by allowing X_t^a -s to be random variables and taking expectations of the two sides of (3) and using the fact that $\mathbb{E}[\min[\cdot]] \leq \min[\mathbb{E}[\cdot]]$. We decompose expectations of incremental sums into incremental sums of conditional expectations and use $\mathbb{E}_t[\cdot]$ to denote expectations conditioned on realization of all random variables up to round t .

Corollary 8. Let X_1^a, X_2^a, \dots for $a \in \{1, \dots, K\}$ be non-negative random variables and let η_t and ρ_t as defined in Lemma 7. Then:

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_a \rho_t(a) X_t^a \right] \right] - \min_a \left(\mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t [X_t^a] \right] \right) \\ & \leq \mathbb{E} \left[\sum_{t=1}^T \frac{\eta_t}{2} \mathbb{E}_t \left[\sum_a \rho_t(a) (X_t^a)^2 \right] \right] + \frac{\ln K}{\eta_T}. \end{aligned}$$

Proof of Theorem 1. We associate X_t^a in (3) with $\tilde{\ell}_t^a$ in the EXP3++ algorithm. We have $\mathbb{E}_t[\tilde{\ell}_t^a] = \ell_t^a$ and since

$$\rho_t(a) = \frac{1}{1 - \sum_{a'} \varepsilon_t(a')} (\tilde{\rho}_t(a) - \varepsilon_t(a)) \geq \tilde{\rho}_t(a) - \varepsilon_t(a)$$

and $\ell_t^a \in [0, 1]$ we also have:

$$\begin{aligned} \mathbb{E}_t \left[\sum_a \rho_t(a) \tilde{\ell}_t^a \right] & \geq \mathbb{E}_t \left[\sum_a (\tilde{\rho}_t(a) - \varepsilon_t(a)) \ell_t^a \right] \\ & \geq \mathbb{E}_t \left[\ell_t^{A_t} \right] - \sum_a \varepsilon_t(a). \end{aligned}$$

As well, we have:

$$\begin{aligned} \mathbb{E}_t \left[\sum_a \rho_t(a) (\tilde{\ell}_t^a)^2 \right] & = \mathbb{E}_t \left[\sum_a \rho_t(a) \left(\frac{\ell_t^{A_t}}{\tilde{\rho}_t(a)} \mathbf{1}_t^a \right)^2 \right] \\ & \leq \mathbb{E}_t \left[\sum_a \frac{\rho_t(a)}{\tilde{\rho}_t(a)^2} \mathbf{1}_t^a \right] = \sum_a \frac{\rho_t(a)}{\tilde{\rho}_t(a)} \\ & = \sum_a \frac{\rho_t(a)}{(1 - \sum_{a'} \varepsilon_t(a')) \rho_t(a) + \varepsilon_t(a)} \leq 2K, \end{aligned}$$

where the last inequality follows by the fact that $(1 - \sum_{a'} \varepsilon_t(a')) \geq \frac{1}{2}$ by the definition of $\varepsilon_t(a)$. Substi-

tution of the above calculations into Corollary 8 yields:

$$\begin{aligned} R(t) & = \mathbb{E} \left[\sum_{t=1}^T \ell_t^{A_t} \right] - \min_a \mathbb{E} \left[\sum_{t=1}^T \ell_t^a \right] \\ & \leq K \sum_{t=1}^T \eta_t + \frac{\ln K}{\eta_T} + \sum_a \varepsilon_t(a) \leq 2K \sum_{t=1}^T \eta_t + \frac{\ln K}{\eta_T}. \end{aligned}$$

The result of the theorem follows by the choice of η_t . \square

The Stochastic Regime

Our proofs are based on the following form of Bernstein's inequality, which is a minor improvement over Cesa-Bianchi & Lugosi (2006, Lemma A.8) based on the ideas from Boucheron et al. (2013, Theorem 2.10).

Theorem 9 (Bernstein's inequality for martingales). Let X_1, \dots, X_n be a martingale difference sequence with respect to filtration $\mathcal{F} = (\mathcal{F}_i)_{1 \leq i \leq n}$ and let $S_i = \sum_{j=1}^i X_j$ be the associated martingale. Assume that there exist positive numbers ν and c , such that $X_j \leq c$ for all j with probability 1 and $\sum_{i=1}^n \mathbb{E}[(X_i)^2 | \mathcal{F}_{i-1}] \leq \nu$ with probability 1. Then for all $b > 0$:

$$\mathbb{P} \left[S_n > \sqrt{2\nu b} + \frac{cb}{3} \right] \leq e^{-b}.$$

We are also using the following technical lemma, which is proved in the supplementary material.

Lemma 10. For any $c > 0$: $\sum_{t=0}^{\infty} e^{-c\sqrt{t}} = O\left(\frac{2}{c^2}\right)$.

The proof of Theorems 2 and 3 is based on the following lemma.

Lemma 11. Let $\{\underline{\varepsilon}_t(a)\}_{t=1}^{\infty}$ be non-increasing deterministic sequences, such that $\underline{\varepsilon}_t(a) \leq \varepsilon_t(a)$ with probability 1 and $\underline{\varepsilon}_t(a) \leq \underline{\varepsilon}_t(a^*)$ for all t and a . Define $\nu_t(a) = \sum_{s=1}^t \frac{1}{\underline{\varepsilon}_s(a)}$ and define the event E_t^a

$$\begin{aligned} & t\Delta(a) - \left(\tilde{L}_t(a) - \tilde{L}_t(a^*) \right) \\ & \leq \sqrt{2(\nu_t(a) + \nu_t(a^*)) b_t} + \frac{1.25b_t}{3\underline{\varepsilon}_t(a^*)}. \end{aligned} \quad (E_t^a)$$

Then for any positive sequence b_1, b_2, \dots and any $t^* \geq 2$ the number of times arm a is played by EXP3++ up to round t is bounded as:

$$\begin{aligned} \mathbb{E}[N_t(a)] & \leq (t^* - 1) + \sum_{s=t^*}^t e^{-b_s} + \sum_{s=t^*}^t \varepsilon_s(a) \mathbf{1}_{\{E_s^a\}} \\ & \quad + \sum_{s=t^*}^t e^{-\eta_s g_{s-1}(a)}, \end{aligned}$$

where

$$g_t(a) = t\Delta(a) - \sqrt{2tb_t \left(\frac{1}{\underline{\varepsilon}_t(a)} + \frac{1}{\underline{\varepsilon}_t(a^*)} \right)} - \frac{1.25b_t}{3\underline{\varepsilon}_t(a^*)}.$$

Proof. Note that elements of the martingale difference sequence $\left\{ \Delta(a) - \left(\tilde{\ell}_t^a - \tilde{\ell}_t^{a^*} \right) \right\}_{t=1}^\infty$ are upper bounded by $\frac{1}{\underline{\varepsilon}_t(a^*)} + 1$. Since $\underline{\varepsilon}_t(a^*) \leq \varepsilon_t(a^*) \leq 1/(2K) \leq 1/4$ we can simplify the upper bound by using $\frac{1}{\underline{\varepsilon}_t(a^*)} + 1 \leq \frac{1.25}{\underline{\varepsilon}_t(a^*)}$. Further note that

$$\begin{aligned} & \sum_{s=1}^t \mathbb{E}_s \left[\left(\Delta(a) - \left(\tilde{\ell}_s^a - \tilde{\ell}_s^{a^*} \right) \right)^2 \right] \\ & \leq \sum_{s=1}^t \mathbb{E}_s \left[\left(\tilde{\ell}_s^a - \tilde{\ell}_s^{a^*} \right)^2 \right] \\ & = \sum_{s=1}^t \left(\mathbb{E}_s \left[\left(\tilde{\ell}_s^a \right)^2 \right] + \mathbb{E}_s \left[\left(\tilde{\ell}_s^{a^*} \right)^2 \right] \right) \\ & \leq \sum_{s=1}^t \left(\frac{1}{\tilde{p}_s(a)} + \frac{1}{\tilde{p}_s(a^*)} \right) \leq \sum_{s=1}^t \left(\frac{1}{\varepsilon_s(a)} + \frac{1}{\varepsilon_s(a^*)} \right) \\ & \leq \sum_{s=1}^t \left(\frac{1}{\underline{\varepsilon}_s(a)} + \frac{1}{\underline{\varepsilon}_s(a^*)} \right) = \nu_t(a) + \nu_t(a^*) \end{aligned}$$

with probability 1. Let \bar{E} denote the complement of event E . Then by Bernstein's inequality $\mathbb{P} \left[\bar{E}_t^a \right] \leq b_t$. The number of times arm a is played up to round t is bounded as:

$$\begin{aligned} \mathbb{E} [N_t(a)] & = \sum_{s=1}^t \mathbb{P} [A_s = a] \\ & = \sum_{s=1}^t \mathbb{P} [A_s = a | E_{s-1}^a] \mathbb{P} [E_{s-1}^a] \\ & \quad + \mathbb{P} [A_s = a | \bar{E}_{s-1}^a] \mathbb{P} [\bar{E}_{s-1}^a] \\ & \leq \sum_{s=1}^t \mathbb{P} [A_s = a | E_{s-1}^a] \mathbb{1}_{\{E_{s-1}^a\}} + \mathbb{P} [\bar{E}_{s-1}^a] \\ & \leq \sum_{s=1}^t \mathbb{P} [A_s = a | E_{s-1}^a] \mathbb{1}_{\{E_{s-1}^a\}} + e^{-b_{s-1}}. \end{aligned}$$

For the terms of the sum above we have:

$$\begin{aligned} \mathbb{P} [A_t = a | E_{t-1}^a] \mathbb{1}_{\{E_{t-1}^a\}} & = \tilde{\rho}_t(a) \mathbb{1}_{\{E_{t-1}^a\}} \\ & \leq (\rho_t(a) + \varepsilon_t(a)) \mathbb{1}_{\{E_{t-1}^a\}} \\ & = \left(\varepsilon_t(a) + \frac{e^{-\eta_t \tilde{L}_{t-1}(a)}}{\sum_{a'} e^{-\eta_t \tilde{L}_{t-1}(a')}} \right) \mathbb{1}_{\{E_{t-1}^a\}} \\ & \leq \left(\varepsilon_t(a) + e^{-\eta_t (\tilde{L}_{t-1}(a) - \tilde{L}_{t-1}(a^*))} \right) \mathbb{1}_{\{E_{t-1}^a\}} \\ & \leq \varepsilon_t(a) \mathbb{1}_{\{E_{t-1}^a\}} + e^{-\eta_t g_{t-1}(a)}, \end{aligned}$$

Where in the last inequality we used the facts that event E_t^a holds and that since $\underline{\varepsilon}_t(a)$ is a non-increasing sequence $\nu_t(a) \leq \frac{t}{\underline{\varepsilon}_t(a)}$. Substitution of this result back into the computation of $\mathbb{E} [N_t(a)]$ completes the proof. \square

Proof of Theorem 2. The proof is based on Lemma 11. Let $b_t = \ln(t\Delta(a)^2)$ and $\underline{\varepsilon}_t(a) = \varepsilon_t(a)$. For any $c \geq 18$ and any $t \geq t^*$, where t^* is the minimal integer for which $t^* \geq \frac{4c^2 K \ln(t^* \Delta(a)^2)^2}{\Delta(a)^4 \ln(K)}$, we have:

$$\begin{aligned} g_t(a) & = t\Delta(a) - \sqrt{2tb_t \left(\frac{1}{\varepsilon_t(a)} + \frac{1}{\varepsilon_t(a^*)} \right)} - \frac{1.25b_t}{3\varepsilon_t(a^*)} \\ & \geq t\Delta(a) - 2\sqrt{\frac{tb_t}{\varepsilon_t(a)}} - \frac{1.25b_t}{3\varepsilon_t(a)} \\ & = t\Delta(a) \left(1 - \frac{2}{\sqrt{c}} - \frac{1.25}{3c} \right) \geq \frac{1}{2}t\Delta(a). \end{aligned}$$

(The choice of t^* ensures that for all suboptimal actions a we have $\varepsilon_t(a) = \xi_t(a)$, which slightly simplifies the calculations. Also note that since $\varepsilon_t(a^*) = \min \left\{ \frac{1}{2K}, \beta_t \right\}$, asymptotically $1/\varepsilon_t(a)$ term in $g_t(a)$ dominates $1/\varepsilon_t(a^*)$ term and with a bit more careful bounding c can be made almost as small as 2.) By substitution of the lower bound on $g_t(a)$ into Lemma 11 we have:

$$\begin{aligned} \mathbb{E} [N_t(a)] & \leq t^* + \frac{\ln(t)}{\Delta(a)^2} + \frac{c \ln(t)^2}{\Delta(a)^2} \\ & \quad + \sum_{s=1}^t \left(e^{-\frac{\Delta(a)}{4}} \sqrt{\frac{(s-1) \ln(K)}{K}} \right) \\ & \leq \frac{c \ln(t)^2}{\Delta(a)^2} + \frac{\ln(t)}{\Delta(a)^2} + O \left(\frac{K}{\Delta(a)^2} \right) + t^*, \end{aligned}$$

where we used Lemma 10 to bound the sum of the exponents. Note that t^* is of order $\tilde{O} \left(\frac{K}{\Delta(a)^4} \right)$. \square

Proof of Theorem 3. Note that since by our definition $\hat{\Delta}_t(a) \leq 1$ the sequence $\underline{\varepsilon}_t(a) = \underline{\varepsilon}_t = \min \left\{ \frac{1}{2K}, \beta_t, \frac{c \ln(t)^2}{t} \right\}$ satisfies the condition of Lemma 11. Also note that for t large enough, so that $t \geq \frac{4c^2 K \ln(t)^4}{\ln K}$, we have $\underline{\varepsilon}_t = \frac{c \ln(t)^2}{t}$. Let $b_t = \ln(t)$ and let t^* be large enough, so that for all $t \geq t^*$ we have $t \geq \frac{4c^2 K \ln(t)^4}{\ln K}$ and $t \geq e^{\frac{1}{\Delta(a)^2}}$. We are going to bound the three terms in the bound on $\mathbb{E} [N_t(a)]$ in Lemma 11. Bounding $\sum_{s=t^*}^t e^{-b_s}$ is easy. For bounding $\sum_{s=t^*}^t \varepsilon_s(a) \mathbb{1}_{\{E_{s-1}^a\}}$ we note that when E_t^a holds and $c \geq 18$ we have:

$$\begin{aligned} \hat{\Delta}_t(a) & \geq \frac{1}{t} \left(\tilde{L}_t(a) - \min_{a'} \tilde{L}_t(a') \right) \geq \frac{1}{t} \left(\tilde{L}_t(a) - \tilde{L}_t(a^*) \right) \\ & \geq \frac{1}{t} g_t(a) = \frac{1}{t} \left(t\Delta(a) - 2\sqrt{\frac{tb_t}{\underline{\varepsilon}_t}} - \frac{1.25b_t}{3\underline{\varepsilon}_t} \right) \quad (4) \\ & = \frac{1}{t} \left(t\Delta(a) - \frac{2t}{\sqrt{c \ln(t)}} - \frac{1.25t}{3c \ln t} \right) \\ & \geq \Delta(a) \left(1 - \frac{2}{\sqrt{c}} - \frac{1.25}{3c} \right) \geq \frac{1}{2}\Delta(a), \end{aligned}$$

where in (4) we used the fact that E_t^a holds and in the last line we used the fact that for $t \geq t^*$ we have $\sqrt{\ln t} \geq 1/\Delta(a)$. Thus

$$\varepsilon_t \mathbb{1}_{\{E_{s-1}^a\}} \leq \frac{c(\ln t)^2}{t\hat{\Delta}_t(a)^2} \leq \frac{4c^2(\ln t)^2}{t\Delta(a)^2}$$

and $\sum_{s=t^*}^t \varepsilon_s \mathbb{1}_{\{E_{s-1}^a\}} = O\left(\frac{\ln(t)^3}{\Delta(a)^2}\right)$. Finally, for the last term in Lemma 11 we have already shown as an intermediate step in the calculation of the bound on $\hat{\Delta}_t(a)$ that for $t \geq t^*$ we have $g_t(a) \geq \frac{1}{2}\Delta(a)$. Therefore, the last term is of order $O\left(\frac{K}{\Delta(a)^2}\right)$. By taking all these calculations together we obtain the result of the theorem. Note that the result holds for any $\eta_t \geq \beta_t$. \square

The Contaminated Stochastic Regime

Proof of Theorem 5. The key element of the previous proof was a high-probability lower bound on $\tilde{L}_t(a) - \tilde{L}_t(a^*)$. We show that we can obtain a similar lower bound in the contaminated setting too. Let $\mathbb{1}_{t,a}^\times$ denote the indicator function of contamination in “location” (t, a) ($\mathbb{1}_{t,a}^\times$ takes value 1 if contamination occurred and 0 otherwise). Let $m_t^a = \mathbb{1}_{t,a}^\times \ell_t^a + (1 - \mathbb{1}_{t,a}^\times) \mu(a)$, in other words, if either a was contaminated on round t then m_t^a is the adversarially assigned value of the loss of arm a on round t and otherwise it is the expected loss. Let $M_t(a) = \sum_{s=1}^t m_s^a$ then $(M_t(a) - M_t(a^*)) - (\tilde{L}_t(a) - \tilde{L}_t(a^*))$ is a martingale. By definition of moderately contaminated after τ rounds process, for $t \geq \tau$ and any suboptimal action a the total number of rounds up to t where either a itself or a^* were contaminated is at most $t\Delta(a)/2$. Therefore, $M_t(a) - M_t(a^*) \geq (t - t\Delta(a)/2) \Delta(a) - t\Delta/2 \geq t\Delta(a)/2$. Define event B_t^a :

$$\frac{t\Delta(a)}{2} - (\tilde{L}_t(a) - \tilde{L}_t(a^*)) \leq 2\sqrt{\nu_t b_t} + \frac{1.25b_t}{3\varepsilon_t}, \quad (B_t^a)$$

where ε_t is defined in the proof of Theorem 3 and $\nu_t = \sum_{s=1}^t \frac{1}{\varepsilon_s}$. Then by Bernstein’s inequality $\mathbb{P}[B_t^a] \leq b_t$. The remainder of the proof is identical to the proof of Theorem 3 with $\Delta(a)$ replaced by $\Delta(a)/2$. \square

The Adversarial Regime with a Gap

The proof of Theorem 6 is based on the following lemma, which is an analogue of Theorems 3 and 5.

Lemma 12. *Under the parametrization given in Theorem 3, the number of times a suboptimal arm a is played by $\text{EXP3++}^{\text{AVG}}$ in an adversarial regime with a gap satisfies:*

$$\mathbb{E}[N_t(a)] \leq \max\left\{t^*, \tau, e^{1/(\Delta(\tau, a)^2)}\right\} + O\left(\frac{\ln(t)^3}{\Delta(\tau, a)^2}\right).$$

Proof. Again, the only modification we need is a high-probability lower bound on $\tilde{L}_t(a) - \tilde{L}_t(a^*)$. We note that $(\lambda_t(a) - \lambda_t(a^*)) - (\tilde{L}_t(a) - \tilde{L}_t(a^*))$ is a martingale and that by definition for $t \geq \tau$ we have $(\lambda_t(a) - \lambda_t(a^*)) \geq t\Delta(\tau, a)$. Define the events W_t^a :

$$t\Delta(\tau, a) - (\tilde{L}_t(a) - \tilde{L}_t(a^*)) \leq 2\sqrt{\nu_t b_t} + \frac{1.25b_t}{3\varepsilon_t}, \quad (W_t^a)$$

where ε_t and ν_t are as in the proof of Theorem 5. By Bernstein’s inequality $\mathbb{P}[W_t^a] \leq b_t$. The remainder of the proof is identical to the proof of Theorem 3. \square

Proof of Theorem 6. Note that by definition $\Delta(\tau, a)$ is a non-decreasing sequence of τ . Since Lemma 12 is a deterministic result it holds for all τ simultaneously and we are free to choose the one that minimizes the bound. \square

5. Empirical Evaluation: Stochastic Regime

We consider the stochastic multiarmed bandit problem with Bernoulli rewards. For all the suboptimal arms the rewards are Bernoulli with bias 0.5 and for the single best arm the reward is Bernoulli with bias $0.5 + \Delta$. We run the experiments with $K = 2$, $K = 10$, and $K = 100$, and $\Delta = 0.1$ and $\Delta = 0.01$ (in total, six combinations of K and Δ). We run each game for 10^7 rounds and make ten repetitions of each experiment. The solid lines in the graphs in Figure 1 represent the mean performance over the experiments and the dashed lines represent the mean plus one standard deviation (std) over the ten repetitions of the corresponding experiment.

In the experiments EXP3++ is parametrized by $\xi_t(a) = \frac{\ln(t\hat{\Delta}_t(a)^2)}{32t\hat{\Delta}_t(a)^2}$, where $\hat{\Delta}_t(a)$ is the empirical estimate of $\Delta(a)$ defined in (2). In order to demonstrate that in the stochastic regime the exploration parameters are in full control of the performance we run the EXP3++ algorithm with two different learning rates. EXP3++^{EMP} corresponds to $\eta_t = \beta_t$ and EXP3++^{ACC} corresponds to $\eta_t = 1$. Note that only the EXP3++^{EMP} has a performance guarantee in the adversarial regime.

We compare EXP3++ algorithm with the EXP3 algorithm (as described in Bubeck & Cesa-Bianchi (2012)), the UCB1 algorithm of Auer et al. (2002a), and Thompson’s sampling. Since it was demonstrated empirically in Seldin et al. (2013) that in the above experiments the performance of Thompson sampling is comparable or superior to the performance of EwS and KL-UCB, the latter two algorithms are excluded from the comparison. For the EXP3++ and the EXP3 algorithms we transform the rewards into losses via $\ell_t^a = 1 - r_t^a$ transformation, other algorithms operate directly on the rewards.

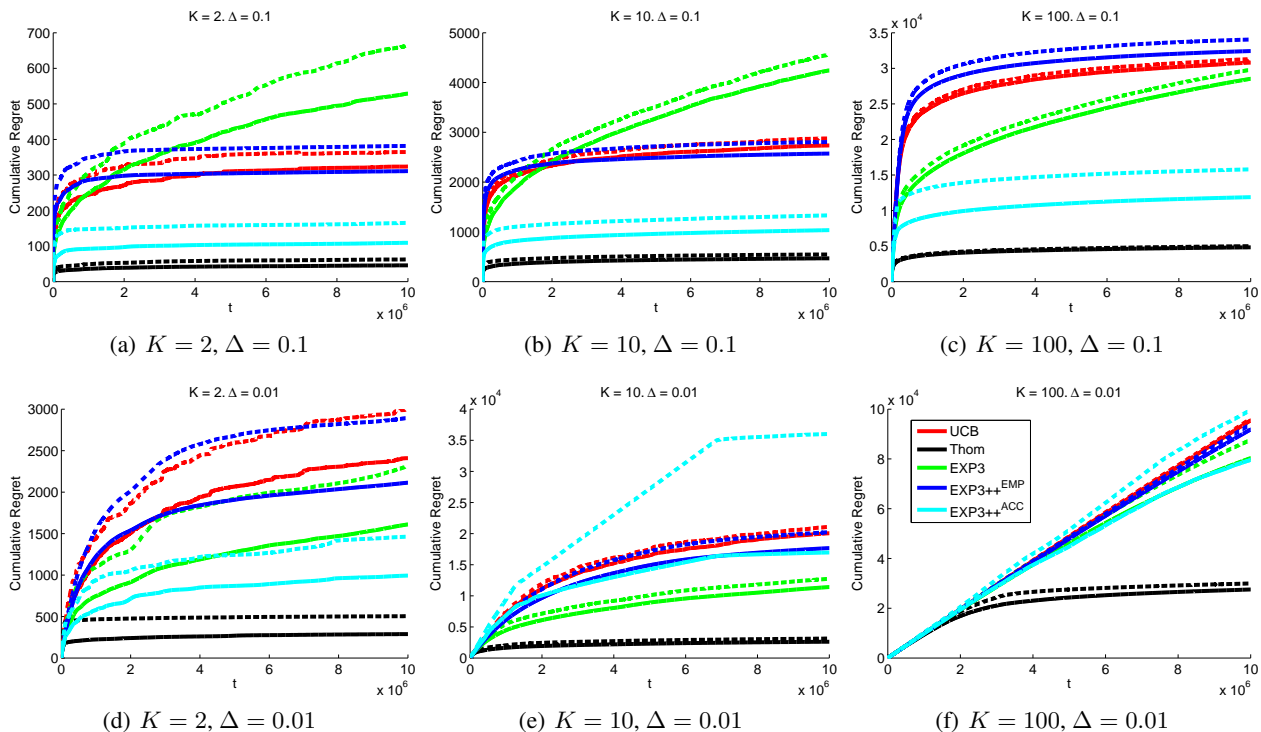


Figure 1. Comparison of UCB1, Thompson sampling (“Thom”), EXP3, and EXP3++ algorithms in the stochastic regime. The legend in figure (f) corresponds to all the figures. EXP3++^{EMP} is the Empirical EXP3++ algorithm and EXP3++^{ACC} is an Accelerated Empirical EXP3++, where we take $\eta_t = 1$. Solid lines correspond to means over 10 repetitions of the corresponding experiments and dashed lines correspond to the means plus one standard deviation.

The results are presented in Figure 1. We see that in all the experiments the performance of EXP3++^{EMP} is almost identical to the performance of UCB1. However, unlike UCB1 and Thompson’s sampling, EXP3++^{EMP} is secured against the possibility that the game is controlled by an adversary. In the supplementary material we show that any deterministic algorithm is vulnerable against an adversary.

The EXP3++^{ACC} algorithm can be seen as a teaser for future work. It performs better than EXP3++^{EMP}, but it does not have the adversarial regime performance guarantee. However, we do not exclude the possibility that by some more sophisticated simultaneous control of η_t and $\varepsilon_t(a)$ -s it may be possible to design an algorithm that will have both better performance in the stochastic regime and a regret guarantee in the adversarial regime. An example of such sophisticated control of the learning rate in the full information games can be found in [de Rooij et al. \(2014\)](#).

6. Discussion

We presented a generalization of the EXP3 algorithm, the EXP3++ algorithm, which augments the EXP3 algorithm with a new control lever in the form exploration parameters $\varepsilon_t(a)$ that are tuned individually for each arm. We have

shown that the new control lever is extremely useful in detecting and exploiting the gap in a wide range of regimes, while the old control lever always keeps the worst-case performance of the algorithm under control. Due to the central role of the EXP3 algorithm in the adversarial analysis that stretches far beyond the adversarial bandits and due to the simplicity of our generalization we believe that our result will lead to a multitude of new algorithms for other problems that exploit the gaps without compromising on the worst-case performance guarantees. There is also room for further improvement of the presented technique that we plan to pursue in future work.

Acknowledgments

The authors would like to thank Sébastien Bubeck and Wouter Koolen for useful discussions and Csaba Szepesvári for bringing up the reference to [Cesa-Bianchi & Fischer \(1998\)](#). This research was supported by an Australian Research Council Australian Laureate Fellowship (FL110100281).

References

- Agrawal, Shipra and Goyal, Navin. Further optimal regret bounds for Thompson sampling. In *AISTATS*, 2013.
- Audibert, Jean-Yves and Bubeck, Sébastien. Minimax policies for adversarial and stochastic bandits. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2009.
- Auer, Peter, Cesa-Bianchi, Nicolò, Freund, Yoav, and Schapire, Robert E. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the Annual Symposium on Foundations of Computer Science*, 1995.
- Auer, Peter, Cesa-Bianchi, Nicolò, and Fischer, Paul. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47, 2002a.
- Auer, Peter, Cesa-Bianchi, Nicolò, Freund, Yoav, and Schapire, Robert E. The nonstochastic multiarmed bandit problem. *SIAM Journal of Computing*, 32(1), 2002b.
- Boucheron, Stéphane, Lugosi, Gábor, and Massart, Pascal. *Concentration Inequalities A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Bubeck, Sébastien. *Bandits Games and Clustering Foundations*. PhD thesis, Université Lille, 2010.
- Bubeck, Sébastien and Cesa-Bianchi, Nicolò. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5, 2012.
- Bubeck, Sébastien and Slivkins, Aleksandrs. The best of both worlds: stochastic and adversarial bandits. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2012.
- Cappé, Olivier, Garivier, Aurélien, Maillard, Odalric-Ambrym, Munos, Rémi, and Stoltz, Gilles. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3), 2013.
- Cesa-Bianchi, Nicolò and Fischer, Paul. Finite-time regret bounds for the multiarmed bandit problem. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1998.
- Cesa-Bianchi, Nicolò and Lugosi, Gábor. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- de Rooij, Steven, van Erven, Tim, Grünwald, Peter D., and Koolen, Wouter M. Follow the leader if you can, hedge if you must. *Journal of Machine Learning Research*, 2014.
- Kaufmann, Emilie, Korda, Nathaniel, and Munos, Rémi. Thompson sampling: An optimal finite time analysis. In *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2012.
- Lai, Tze Leung and Robbins, Herbert. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6, 1985.
- Maillard, Odalric-Ambrym. *Apprentissage Séquentiel: Bandits, Statistique et Renforcement*. PhD thesis, INRIA Lille, 2011.
- Robbins, Herbert. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 1952.
- Seldin, Yevgeny, Szepesvári, Csaba, Auer, Peter, and Abbasi-Yadkori, Yasin. Evaluation and analysis of the performance of the EXP3 algorithm in stochastic environments. In *JMLR Workshop and Conference Proceedings*, volume 24 (EWRL), 2013.
- Stoltz, Gilles. *Incomplete Information and Internal Regret in Prediction of Individual Sequences*. PhD thesis, Université Paris-Sud, 2005.
- Thompson, William R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25, 1933.