

Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems

By Sébastien Bubeck and Nicolò Cesa-Bianchi

Contents

1	Introduction	2
2	Stochastic Bandits: Fundamental Results	9
2.1	Optimism in Face of Uncertainty	10
2.2	Upper Confidence Bound (UCB) Strategies	11
2.3	Lower Bound	13
2.4	Refinements and Bibliographic Remarks	17
3	Adversarial Bandits: Fundamental Results	22
3.1	Pseudo-regret Bounds	23
3.2	High Probability and Expected Regret Bounds	28
3.3	Lower Bound	33
3.4	Refinements and Bibliographic Remarks	37
4	Contextual Bandits	43
4.1	Bandits with Side Information	44
4.2	The Expert Case	45
4.3	Stochastic Contextual Bandits	53
4.4	The Multiclass Case	56
4.5	Bibliographic Remarks	62

5 Linear Bandits	64
5.1 Exp2 (Expanded Exp) with John's Exploration	65
5.2 Online Mirror Descent (OMD)	69
5.3 Online Stochastic Mirror Descent (OSMD)	74
5.4 Online Combinatorial Optimization	76
5.5 Improved Regret Bounds for Bandit Feedback	81
5.6 Refinements and Bibliographic Remarks	84
6 Nonlinear Bandits	88
6.1 Two-Point Bandit Feedback	89
6.2 One-Point Bandit Feedback	94
6.3 Nonlinear Stochastic Bandits	96
6.4 Bibliographic Remarks	101
7 Variants	103
7.1 Markov Decision Processes, Restless and Sleeping Bandits	104
7.2 Pure Exploration Problems	105
7.3 Dueling Bandits	107
7.4 Discovery with Probabilistic Expert Advice	107
7.5 Many-Armed Bandits	108
7.6 Truthful Bandits	109
7.7 Concluding Remarks	109
Acknowledgments	111
References	112

Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems

Sébastien Bubeck¹ and Nicolò Cesa-Bianchi²

¹ *Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA, sbubeck@princeton.edu*

² *Dipartimento di Informatica, Università degli Studi di Milano, Milano 20135, Italy, nicolo.cesa-bianchi@unimi.it*

Abstract

Multi-armed bandit problems are the most basic examples of sequential decision problems with an exploration–exploitation trade-off. This is the balance between staying with the option that gave highest payoffs in the past and exploring new options that might give higher payoffs in the future. Although the study of bandit problems dates back to the 1930s, exploration–exploitation trade-offs arise in several modern applications, such as ad placement, website optimization, and packet routing. Mathematically, a multi-armed bandit is defined by the payoff process associated with each option. In this monograph, we focus on two extreme cases in which the analysis of regret is particularly simple and elegant: i.i.d. payoffs and adversarial payoffs. Besides the basic setting of finitely many actions, we also analyze some of the most important variants and extensions, such as the contextual bandit model.

1

Introduction

A multi-armed bandit problem (or, simply, a bandit problem) is a sequential allocation problem defined by a set of actions. At each time step, a unit resource is allocated to an action and some observable payoff is obtained. The goal is to maximize the total payoff obtained in a sequence of allocations. The name *bandit* refers to the colloquial term for a slot machine (“one-armed bandit” in American slang). In a casino, a sequential allocation problem is obtained when the player is facing many slot machines at once (a “multi-armed bandit”) and must repeatedly choose where to insert the next coin.

Bandit problems are basic instances of sequential decision making with limited information and naturally address the fundamental trade-off between exploration and exploitation in sequential experiments. Indeed, the player must balance the exploitation of actions that did well in the past and the exploration of actions that might give higher payoffs in the future.

Although the original motivation of Thompson [162] for studying bandit problems came from clinical trials (when different treatments are available for a certain disease and one must decide which treatment to use on the next patient), modern technologies have created

many opportunities for new applications, and bandit problems now play an important role in several industrial domains. In particular, online services are natural targets for bandit algorithms, because there one can benefit from adapting the service to the individual sequence of requests. We now describe a few concrete examples in various domains.

Ad placement is the problem of deciding which advertisement to display on the web page delivered to the next visitor of a website. Similarly, website optimization deals with the problem of sequentially choosing design elements (font, images, layout) for the web page. Here the payoff is associated with visitor’s actions, e.g., clickthroughs or other desired behaviors. Of course there are important differences with the basic bandit problem: in ad placement the pool of available ads (bandit arms) may change over time, and there might be a limit on the number of times each ad could be displayed.

In source routing a sequence of packets must be routed from a source host to a destination host in a given network, and the protocol allows to choose a specific source-destination path for each packet to be sent. The (negative) payoff is the time it takes to deliver a packet, and depends additively on the congestion of the edges in the chosen path.

In computer game-playing, each move is chosen by simulating and evaluating many possible game continuations after the move. Algorithms for bandits (more specifically, for a tree-based version of the bandit problem) can be used to explore more efficiently the huge tree of game continuations by focusing on the most promising subtrees. This idea has been successfully implemented in the MoGo player of Gelly et al. [85], which plays Go at world-class level. MoGo is based on the UCT strategy for hierarchical bandits of Kocsis and Szepesvári [123], which in turn is derived from the UCB bandit algorithm — see Section 2.

There are three fundamental formalizations of the bandit problem depending on the assumed nature of the reward process: stochastic, adversarial, and Markovian. Three distinct playing strategies have been shown to effectively address each specific bandit model: the UCB algorithm in the stochastic case, the Exp3 randomized algorithm in the adversarial case, and the so-called Gittins indices in the Markovian case. In this monograph, we focus on stochastic and adversarial bandits,

4 Introduction

and refer the reader to the monograph by Mahajan and Teneketzis [130] or to the recent monograph by Gittins et al. [86] for an extensive analysis of Markovian bandits.

In order to analyze the behavior of a player or forecaster (i.e., the agent implementing a bandit strategy), we may compare its performance with that of an optimal strategy that, for any horizon of n time steps, consistently plays the arm that is best in the first n steps. In other terms, we may study the *regret* of the forecaster for not playing always optimally. More specifically, given $K \geq 2$ arms and sequences $X_{i,1}, X_{i,2}, \dots$ of unknown rewards associated with each arm $i = 1, \dots, K$, we study forecasters that at each time step $t = 1, 2, \dots$ select an arm I_t and receive the associated reward $X_{I_t,t}$. The regret after n plays I_1, \dots, I_n is defined by

$$R_n = \max_{i=1, \dots, K} \sum_{t=1}^n X_{i,t} - \sum_{t=1}^n X_{I_t,t}. \quad (1.1)$$

If the time horizon is not known in advance we say that the forecaster is *anytime*.

In general, both rewards $X_{i,t}$ and forecaster's choices I_t might be stochastic. This allows to distinguish between the two following notions of averaged regret: the *expected regret*

$$\mathbb{E} R_n = \mathbb{E} \left[\max_{i=1, \dots, K} \sum_{t=1}^n X_{i,t} - \sum_{t=1}^n X_{I_t,t} \right] \quad (1.2)$$

and the *pseudo-regret*

$$\bar{R}_n = \max_{i=1, \dots, K} \mathbb{E} \left[\sum_{t=1}^n X_{i,t} - \sum_{t=1}^n X_{I_t,t} \right]. \quad (1.3)$$

In both definitions, the expectation is taken with respect to the random draw of both rewards and forecaster's actions. Note that pseudo-regret is a weaker notion of regret, since one competes against the action which is optimal only in expectation. The expected regret, instead, is the expectation of the regret with respect to the action which is optimal on the sequence of reward realizations. More formally one has $\bar{R}_n \leq \mathbb{E} R_n$.

In the original formalization of Robbins [146], which builds on the work of Wald [164] — see also Arrow et al. [16], each arm $i = 1, \dots, K$

corresponds to an unknown probability distribution ν_i on $[0, 1]$, and rewards $X_{i,t}$ are independent draws from the distribution ν_i corresponding to the selected arm.

The stochastic bandit problem

Known parameters: number of arms K and (possibly) number of rounds $n \geq K$.

Unknown parameters: K probability distributions ν_1, \dots, ν_K on $[0, 1]$.

For each round $t = 1, 2, \dots$

- (1) the forecaster chooses $I_t \in \{1, \dots, K\}$;
- (2) given I_t , the environment draws the reward $X_{I_t,t} \sim \nu_{I_t}$ independently from the past and reveals it to the forecaster.

For $i = 1, \dots, K$ we denote by μ_i the mean of ν_i (mean reward of arm i). Let

$$\mu^* = \max_{i=1, \dots, K} \mu_i \quad \text{and} \quad i^* \in \operatorname{argmax}_{i=1, \dots, K} \mu_i.$$

In the stochastic setting, it is easy to see that the pseudo-regret can be written as

$$\bar{R}_n = n\mu^* - \sum_{t=1}^n \mathbb{E}[\mu_{I_t}]. \quad (1.4)$$

The analysis of the stochastic bandit model was pioneered in the seminal paper of Lai and Robbins [125], who introduced the technique of upper confidence bounds for the asymptotic analysis of regret. In Section 2 we describe this technique using the simpler formulation of Agrawal [9], which naturally lends itself to a finite-time analysis.

In parallel to the research on stochastic bandits, a game-theoretic formulation of the trade-off between exploration and exploitation has been independently investigated, although for quite some time this alternative formulation was not recognized as an instance of the multi-armed bandit problem. In order to motivate these game-theoretic bandits, consider again the initial example of gambling on slot machines. We now assume that we are in a rigged casino, where for each slot machine $i = 1, \dots, K$ and time step $t \geq 1$ the owner sets the gain $X_{i,t}$ to some arbitrary (and possibly maliciously chosen) value

$g_{i,t} \in [0, 1]$. Note that it is not in the interest of the owner to simply set all the gains to zero (otherwise, no gamblers would go to that casino). Now recall that a forecaster selects sequentially one arm $I_t \in \{1, \dots, K\}$ at each time step $t = 1, 2, \dots$ and observes (and earns) the gain $g_{I_t,t}$. Is it still possible to minimize regret in such a setting?

Following a standard terminology, we call adversary, or opponent, the mechanism setting the sequence of gains for each arm. If this mechanism is independent of the forecaster's actions, then we call it an *oblivious* adversary. In general, however, the adversary may adapt to the forecaster's past behavior, in which case we speak of a *nonoblivious* adversary. For instance, in the rigged casino the owner may observe the way a gambler plays in order to design even more evil sequences of gains. Clearly, the distinction between oblivious and nonoblivious adversary is only meaningful when the player is randomized (if the player is deterministic, then the adversary can pick a bad sequence of gains right at the beginning of the game by simulating the player's future actions). Note, however, that in the presence of a nonoblivious adversary the interpretation of regret is ambiguous. Indeed, in this case the assignment of gains $g_{i,t}$ to arms $i = 1, \dots, K$ made by the adversary at each step t is allowed to depend on the player's past randomized actions I_1, \dots, I_{t-1} . In other words, $g_{i,t} = g_{i,t}(I_1, \dots, I_{t-1})$ for each i and t . Now, the regret compares the player's cumulative gain to that obtained by playing the single best arm for the first n rounds. However, had the player consistently chosen the same arm i in each round, namely $I_t = i$ for $t = 1, \dots, n$, the adversarial gains $g_{i,t}(I_1, \dots, I_{t-1})$ would have been possibly different than those actually experienced by the player.

The study of nonoblivious regret is mainly motivated by the connection between regret minimization and equilibria in games — see, e.g., [24, Section 9]. Here we just observe that game-theoretic equilibria are indeed defined similarly to regret: in equilibrium, the player has no incentive to behave differently, provided the opponent does not react to changes in the player's behavior. Interestingly, regret minimization has been also studied against *reactive opponents*, see for instance the works of Pucci de Farias and Megiddo [144] and Arora et al. [14].

The adversarial bandit problem

Known parameters: number of arms $K \geq 2$ and (possibly) number of rounds $n \geq K$.

For each round $t = 1, 2, \dots$

- (1) the forecaster chooses $I_t \in \{1, \dots, K\}$, possibly with the help of external randomization,
- (2) simultaneously, the adversary selects a gain vector $g_t = (g_{1,t}, \dots, g_{K,t}) \in [0, 1]^K$, possibly with the help of external randomization, and
- (3) the forecaster receives (and observes) the reward $g_{I_t,t}$, while the gains of the other arms are not observed.

In this adversarial setting the goal is to obtain regret bounds in high probability or in expectation with respect to any possible randomization in the strategies used by the forecaster or the opponent, and irrespective of the opponent. In the case of a nonoblivious adversary this is not an easy task, and for this reason we usually start by bounding the pseudo-regret

$$\bar{R}_n = \max_{i=1, \dots, K} \mathbb{E} \left[\sum_{t=1}^n g_{i,t} - \sum_{t=1}^n g_{I_t,t} \right].$$

Note that the randomization of the adversary is not very important here since we ask for bounds which hold for any opponent. On the other hand, it is fundamental to allow randomization for the forecaster — see Section 3 for details and basic results in the adversarial bandit model. This adversarial, or nonstochastic, version of the bandit problem was originally proposed as a way of playing an unknown game against an opponent. The problem of playing a game repeatedly, now a classical topic in game theory, was initiated by the groundbreaking work of James Hannan and David Blackwell. In Hannan’s seminal paper Hannan [92], the game (i.e., the payoff matrix) is assumed to be known by the player, who also observes the opponent’s moves in each play. Later, Baños [28] considered the problem of a repeated unknown game, where in each game round the player only observes its own payoff. This problem turns out to be exactly equivalent to the adversarial bandit

problem with a nonoblivious adversary. Simpler strategies for playing unknown games were more recently proposed by Foster and Vohra [81] and Hart and Mas-Colell [93, 94]. Approximately at the same time, the problem was re-discovered in computer science by Auer et al. [24]. It was them who made apparent the connection to stochastic bandits by coining the term nonstochastic multi-armed bandit problem.

The third fundamental model of multi-armed bandits assumes that the reward processes are neither i.i.d. (like in stochastic bandits) nor adversarial. More precisely, arms are associated with K Markov processes, each with its own state space. Each time an arm i is chosen in state s , a stochastic reward is drawn from a probability distribution $\nu_{i,s}$, and the state of the reward process for arm i changes in a Markovian fashion, based on an underlying stochastic transition matrix M_i . Both reward and new state are revealed to the player. On the other hand, the state of arms that are not chosen remains unchanged. Going back to our initial interpretation of bandits as sequential resource allocation processes, here we may think of K competing projects that are sequentially allocated a unit resource of work. However, unlike the previous bandit models, in this case the state of a project that gets the resource may change. Moreover, the underlying stochastic transition matrices M_i are typically assumed to be known, thus the optimal policy can be computed via dynamic programming and the problem is essentially of computational nature. The seminal result of Gittins [87] provides an optimal greedy policy which can be computed efficiently.

A notable special case of Markovian bandits is that of Bayesian bandits. These are parametric stochastic bandits, where the parameters of the reward distributions are assumed to be drawn from known priors, and the regret is computed by also averaging over the draw of parameters from the prior. The Markovian state change associated with the selection of an arm corresponds here to updating the posterior distribution of rewards for that arm after observing a new reward.

Markovian bandits are a standard model in the areas of Operations Research and Economics. However, the techniques used in their analysis are significantly different from those used to analyze stochastic and adversarial bandits. For this reason, in this monograph we do not cover Markovian bandits and their many variants.

2

Stochastic Bandits: Fundamental Results

We start by recalling the basic definitions for the stochastic bandit problem. Each arm $i \in \{1, \dots, K\}$ corresponds to an unknown probability distribution ν_i . At each time step $t = 1, 2, \dots$, the forecaster selects some arm $I_t \in \{1, \dots, K\}$ and receives a reward $X_{I_t, t}$ drawn from ν_{I_t} (independently from the past). Denote by μ_i the mean of arm i and define

$$\mu^* = \max_{i=1, \dots, K} \mu_i \quad \text{and} \quad i^* \in \operatorname{argmax}_{i=1, \dots, K} \mu_i.$$

We focus here on the pseudo-regret, which is defined as

$$\bar{R}_n = n\mu^* - \mathbb{E} \sum_{t=1}^n \mu_{I_t}. \quad (2.1)$$

We choose the pseudo-regret as our main quantity of interest because in a stochastic framework it is more natural to compete against the optimal action in expectation, rather than the optimal action on the sequence of realized rewards (as in the definition of the plain regret in Equation (1.1)). Furthermore, because of the order of magnitude of typical random fluctuations, in general one cannot hope to prove a bound on the expected regret in Equation (1.2) better than $\Theta(\sqrt{n})$.

On the contrary, the pseudo-regret can be controlled so well that we are able to bound it by a *logarithmic* function of n .

In the following we also use a different formula for the pseudo-regret. Let $T_i(s) = \sum_{t=1}^s \mathbb{1}_{I_t=i}$ denote the number of times the player selected arm i on the first s rounds. Let $\Delta_i = \mu^* - \mu_i$ be the suboptimality parameter of arm i . Then the pseudo-regret can be written as:

$$\bar{R}_n = \left(\sum_{i=1}^K \mathbb{E} T_i(n) \right) \mu^* - \mathbb{E} \sum_{i=1}^K T_i(n) \mu_i = \sum_{i=1}^K \Delta_i \mathbb{E} T_i(n).$$

2.1 Optimism in Face of Uncertainty

The difficulty of the stochastic multi-armed bandit problem lies in the *exploration–exploitation dilemma* that the forecaster is facing. Indeed, there is an intrinsic trade-off between *exploiting* the current knowledge to focus on the arm that seems to yield the highest rewards and *exploring* further the other arms to identify with better precision which arm is actually the best. As we shall see, the key to obtain a good strategy for this problem is, in a certain sense, to simultaneously perform exploration and exploitation.

A simple heuristic principle for doing that is the so-called *optimism in face of uncertainty*. The idea is very general and applies to many sequential decision-making problems in uncertain environments. Assume that the forecaster has accumulated some data on the environment and must decide how to act next. First, a set of “plausible” environments which are “consistent” with the data (typically, through concentration inequalities) is constructed. Then, the most “favorable” environment is identified in this set. Based on that, the heuristic prescribes that the decision which is optimal in this most favorable and plausible environment should be made. As we see below, this principle gives simple and yet almost optimal algorithms for the stochastic multi-armed bandit problem. More complex algorithms for various extensions of the stochastic multi-armed bandit problem are also based on the same idea which, along with the exponential weighting scheme presented in Section 3, is an algorithmic cornerstone of regret analysis in bandits.

2.2 Upper Confidence Bound (UCB) Strategies

In this section we assume that the distribution of rewards X satisfies the following moment conditions. There exists a convex function¹ ψ on the reals such that, for all $\lambda \geq 0$,

$$\ln \mathbb{E} e^{\lambda(X - \mathbb{E}[X])} \leq \psi(\lambda) \quad \text{and} \quad \ln \mathbb{E} e^{\lambda(\mathbb{E}[X] - X)} \leq \psi(\lambda). \quad (2.2)$$

For example, when $X \in [0, 1]$ one can take $\psi(\lambda) = \frac{\lambda^2}{8}$. In this case Equation (2.2) is known as Hoeffding's lemma.

We attack the stochastic multi-armed bandit using the optimism in face of uncertainty principle. In order to do so, we use assumption (Equation (2.2)) to construct an upper bound estimate on the mean of each arm at some fixed confidence level and then choose the arm that looks best under this estimate. We need a standard notion from convex analysis: the Legendre–Fenchel transform of ψ , defined by

$$\psi^*(\varepsilon) = \sup_{\lambda \in \mathbb{R}} (\lambda \varepsilon - \psi(\lambda)).$$

For instance, if $\psi(x) = e^x$ then $\psi^*(x) = x \ln x - x$ for $x > 0$. If $\psi(x) = \frac{1}{p}|x|^p$ then $\psi^*(x) = \frac{1}{q}|x|^q$ for any pair $1 < p, q < \infty$ such that $\frac{1}{p} + \frac{1}{q} = 1$ — see also Section 5.2, where the same notion is used in a different bandit model.

Let $\hat{\mu}_{i,s}$ be the sample mean of rewards obtained by pulling arm i for s times. Note that since the rewards are i.i.d., we have that in distribution $\hat{\mu}_{i,s}$ is equal to $\frac{1}{s} \sum_{t=1}^s X_{i,t}$.

Using Markov's inequality, from Equation (2.2) one obtains that

$$\mathbb{P}(\mu_i - \hat{\mu}_{i,s} > \varepsilon) \leq e^{-s\psi^*(\varepsilon)}. \quad (2.3)$$

In other words, with probability at least $1 - \delta$,

$$\hat{\mu}_{i,s} + (\psi^*)^{-1} \left(\frac{1}{s} \ln \frac{1}{\delta} \right) > \mu_i.$$

We thus consider the following strategy, called (α, ψ) -UCB, where $\alpha > 0$ is an input parameter: At time t , select

$$I_t \in \operatorname{argmax}_{i=1, \dots, K} \left[\hat{\mu}_{i, T_i(t-1)} + (\psi^*)^{-1} \left(\frac{\alpha \ln t}{T_i(t-1)} \right) \right].$$

We can prove the following simple bound.

¹One can easily generalize the discussion to functions ψ defined only on an interval $[0, b]$.

Theorem 2.1 (Pseudo-regret of (α, ψ) -UCB). Assume that the reward distributions satisfy Equation (2.2). Then (α, ψ) -UCB with $\alpha > 2$ satisfies

$$\bar{R}_n \leq \sum_{i: \Delta_i > 0} \left(\frac{\alpha \Delta_i}{\psi^*(\Delta_i/2)} \ln n + \frac{\alpha}{\alpha - 2} \right).$$

In case of $[0, 1]$ -valued random variables, taking $\psi(\lambda) = \frac{\lambda^2}{8}$ in Equation (2.2) — the Hoeffding's Lemma — gives $\psi^*(\varepsilon) = 2\varepsilon^2$, which in turns gives the following pseudo-regret bound

$$\bar{R}_n \leq \sum_{i: \Delta_i > 0} \left(\frac{2\alpha}{\Delta_i} \ln n + \frac{\alpha}{\alpha - 2} \right). \quad (2.4)$$

In this important special case of bounded random variables we refer to (α, ψ) -UCB simply as α -UCB.

Proof. First note that if $I_t = i$, then at least one of the three following equations must be true:

$$\hat{\mu}_{i^*, T_{i^*}(t-1)} + (\psi^*)^{-1} \left(\frac{\alpha \ln t}{T_{i^*}(t-1)} \right) \leq \mu^* \quad (2.5)$$

$$\hat{\mu}_{i, T_i(t-1)} > \mu_i + (\psi^*)^{-1} \left(\frac{\alpha \ln t}{T_i(t-1)} \right) \quad (2.6)$$

$$T_i(t-1) < \frac{\alpha \ln n}{\psi^*(\Delta_i/2)}. \quad (2.7)$$

Indeed, assume that the three equations are all false, then we have:

$$\begin{aligned} \hat{\mu}_{i^*, T_{i^*}(t-1)} + (\psi^*)^{-1} \left(\frac{\alpha \ln t}{T_{i^*}(t-1)} \right) &> \mu^* \\ &= \mu_i + \Delta_i \\ &\geq \mu_i + 2(\psi^*)^{-1} \left(\frac{\alpha \ln t}{T_i(t-1)} \right) \\ &\geq \hat{\mu}_{i, T_i(t-1)} + (\psi^*)^{-1} \left(\frac{\alpha \ln t}{T_i(t-1)} \right) \end{aligned}$$

which implies, in particular, that $I_t \neq i$. In other words, letting

$$u = \left\lceil \frac{\alpha \ln n}{\psi^*(\Delta_i/2)} \right\rceil$$

we just proved

$$\begin{aligned} \mathbb{E}T_i(n) &= \mathbb{E} \sum_{t=1}^n \mathbb{1}_{I_t=i} \leq u + \mathbb{E} \sum_{t=u+1}^n \mathbb{1}_{I_t=i} \text{ and (2.7) is false} \\ &\leq u + \mathbb{E} \sum_{t=u+1}^n \mathbb{1}_{(2.5) \text{ or } (2.6) \text{ is true}} \\ &= u + \sum_{t=u+1}^n \mathbb{P}((2.5) \text{ is true}) + \mathbb{P}((2.6) \text{ is true}). \end{aligned}$$

Thus it suffices to bound the probability of the events (Equations (2.5) and (2.6)). Using a union bound and Equation (2.3) one directly obtains:

$$\begin{aligned} \mathbb{P}((2.5) \text{ is true}) &\leq \mathbb{P}\left(\exists s \in \{1, \dots, t\} : \widehat{\mu}_{i^*,s} + (\psi^*)^{-1}\left(\frac{\alpha \ln t}{s}\right) \leq \mu^*\right) \\ &\leq \sum_{s=1}^t \mathbb{P}\left(\widehat{\mu}_{i^*,s} + (\psi^*)^{-1}\left(\frac{\alpha \ln t}{s}\right) \leq \mu^*\right) \\ &\leq \sum_{s=1}^t \frac{1}{t^\alpha} = \frac{1}{t^{\alpha-1}}. \end{aligned}$$

The same upper bound holds for Equation (2.6). Straightforward computations conclude the proof. \square

2.3 Lower Bound

We now show that the result of the previous section is essentially unimprovable when the reward distributions are Bernoulli. For $p, q \in [0, 1]$ we denote by $\text{kl}(p, q)$ the Kullback–Leibler divergence between a Bernoulli of parameter p and a Bernoulli of parameter q , defined as

$$\text{kl}(p, q) = p \ln \frac{p}{q} + (1 - p) \ln \frac{1 - p}{1 - q}.$$

Theorem 2.2 (Distribution-dependent lower bound). Consider a strategy that satisfies $\mathbb{E}T_i(n) = o(n^a)$ for any set of Bernoulli reward distributions, any arm i with $\Delta_i > 0$, and any $a > 0$. Then, for any set of Bernoulli reward distributions the following holds

$$\liminf_{n \rightarrow +\infty} \frac{\bar{R}_n}{\ln n} \geq \sum_{i: \Delta_i > 0} \frac{\Delta_i}{\text{kl}(\mu_i, \mu^*)}.$$

In order to compare this result with Equation (2.4) we use the following standard inequalities (the left-hand side follows from Pinsker's inequality and the right-hand side simply uses $\ln x \leq x - 1$),

$$2(p - q)^2 \leq \text{kl}(p, q) \leq \frac{(p - q)^2}{q(1 - q)}. \quad (2.8)$$

Proof. The proof is organized in three steps. For simplicity, we only consider the case of two arms.

First step: Notations

Without loss of generality assume that arm 1 is optimal and arm 2 is suboptimal, that is, $\mu_2 < \mu_1 < 1$. Let $\varepsilon > 0$. Since $x \mapsto \text{kl}(\mu_2, x)$ is continuous one can find $\mu'_2 \in (\mu_1, 1)$ such that

$$\text{kl}(\mu_2, \mu'_2) \leq (1 + \varepsilon)\text{kl}(\mu_2, \mu_1). \quad (2.9)$$

We use the notation \mathbb{E}', \mathbb{P}' when we integrate with respect to the modified bandit where the parameter of arm 2 is replaced by μ'_2 . We want to compare the behavior of the forecaster on the initial and modified bandits. In particular, we prove that with a big enough probability the forecaster cannot distinguish between the two problems. Then, using the fact that we have a good forecaster by hypothesis, we know that the algorithm does not make too many mistakes on the modified bandit where arm 2 is optimal. In other words, we have a lower bound on the number of times the optimal arm is played. This reasoning implies a lower bound on the number of times arm 2 is played in the initial problem.

We now slightly change the notation for rewards and denote by $X_{2,1}, \dots, X_{2,n}$ the sequence of random variables obtained when pulling arm 2 for n times (that is, $X_{2,s}$ is the reward obtained from the s -th pull). For $s \in \{1, \dots, n\}$, let

$$\widehat{\text{kl}}_s = \sum_{t=1}^s \ln \frac{\mu_2 X_{2,t} + (1 - \mu_2)(1 - X_{2,t})}{\mu'_2 X_{2,t} + (1 - \mu'_2)(1 - X_{2,t})}.$$

Note that, with respect to the initial bandit, $\widehat{\text{kl}}_{T_2(n)}$ is the (nonre-normalized) empirical estimate of $\text{kl}(\mu_2, \mu'_2)$ at time n , since in that case the process $(X_{2,s})$ is i.i.d. from a Bernoulli of parameter μ_2 . Another important property is the following: for any event A in the σ -algebra generated by $X_{2,1}, \dots, X_{2,n}$ the following change-of-measure identity holds:

$$\mathbb{P}'(A) = \mathbb{E}[\mathbb{1}_A \exp(-\widehat{\text{kl}}_{T_2(n)})]. \quad (2.10)$$

In order to link the behavior of the forecaster on the initial and modified bandits we introduce the event

$$C_n = \left\{ T_2(n) < \frac{1 - \varepsilon}{\text{kl}(\mu_2, \mu'_2)} \ln(n) \text{ and } \widehat{\text{kl}}_{T_2(n)} \leq \left(1 - \frac{\varepsilon}{2}\right) \ln(n) \right\}. \quad (2.11)$$

Second step: $\mathbb{P}(C_n) = o(1)$

By Equations (2.10) and (2.11) one has

$$\mathbb{P}'(C_n) = \mathbb{E} \mathbb{1}_{C_n} \exp(-\widehat{\text{kl}}_{T_2(n)}) \geq e^{-(1-\varepsilon/2)\ln(n)} \mathbb{P}(C_n).$$

Introduce the shorthand

$$f_n = \frac{1 - \varepsilon}{\text{kl}(\mu_2, \mu'_2)} \ln(n).$$

Using again Equation (2.11) and Markov's inequality, the above implies

$$\begin{aligned} \mathbb{P}(C_n) &\leq n^{(1-\varepsilon/2)} \mathbb{P}'(C_n) \leq n^{(1-\varepsilon/2)} \mathbb{P}'(T_2(n) < f_n) \\ &\leq n^{(1-\varepsilon/2)} \frac{\mathbb{E}'[n - T_2(n)]}{n - f_n}. \end{aligned}$$

Now note that in the modified bandit arm 2 is the unique optimal arm. Hence the assumption that for any bandit, any suboptimal arm i , and any $a > 0$, the strategy satisfies $\mathbb{E}T_i(n) = o(n^a)$, implies that

$$\mathbb{P}(C_n) \leq n^{(1-\varepsilon/2)} \frac{\mathbb{E}'[n - T_2(n)]}{n - f_n} = o(1).$$

Third step: $\mathbb{P}(T_2(n) < f_n) = o(1)$.

Observe that

$$\begin{aligned} \mathbb{P}(C_n) &\geq \mathbb{P}\left(T_2(n) < f_n \text{ and } \max_{s \leq f_n} \widehat{\text{kl}}_s \leq \left(1 - \frac{\varepsilon}{2}\right) \ln(n)\right) \\ &= \mathbb{P}\left(T_2(n) < f_n \text{ and } \frac{\text{kl}(\mu_2, \mu'_2)}{(1 - \varepsilon) \ln(n)} \right. \\ &\quad \left. \times \max_{s \leq f_n} \widehat{\text{kl}}_s \leq \frac{1 - \varepsilon/2}{1 - \varepsilon} \text{kl}(\mu_2, \mu'_2)\right). \end{aligned} \quad (2.12)$$

Now we use the maximal version of the strong law of large numbers: for any sequence (X_t) of independent real random variables with positive mean $\mu > 0$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n X_t = \mu \quad a.s. \quad \text{implies} \quad \lim_{n \rightarrow \infty} \frac{1}{n} \max_{s=1, \dots, n} \sum_{t=1}^s X_t = \mu \quad a.s.$$

See, e.g., [43, Lemma 10.5].

Since $\text{kl}(\mu_2, \mu'_2) > 0$ and $\frac{1-\varepsilon/2}{1-\varepsilon} > 1$, we deduce that

$$\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{\text{kl}(\mu_2, \mu'_2)}{(1 - \varepsilon) \ln(n)} \times \max_{s \leq f_n} \widehat{\text{kl}}_s \leq \frac{1 - \varepsilon/2}{1 - \varepsilon} \text{kl}(\mu_2, \mu'_2)\right) = 1.$$

Thus, by the result of the second step and Equation (2.12), we get

$$\mathbb{P}(T_2(n) < f_n) = o(1).$$

Now recalling that $f_n = \frac{1-\varepsilon}{\text{kl}(\mu_2, \mu'_2)} \ln(n)$, and using Equation (2.9), we obtain

$$\mathbb{E}T_2(n) \geq (1 + o(1)) \frac{1 - \varepsilon}{1 + \varepsilon} \frac{\ln(n)}{\text{kl}(\mu_2, \mu_1)}$$

which concludes the proof. \square

2.4 Refinements and Bibliographic Remarks

The UCB strategy presented in Section 2.2 was introduced by Auer et al. [23] for bounded random variables. Theorem 2.2 is extracted from Lai and Robbins [125]. Note that in this last paper the result is more general than ours, which is restricted to Bernoulli distributions. Although Burnetas and Katchakis [56] prove an even more general lower bound, Theorem 2.2 and the UCB regret bound provide a reasonably complete solution to the problem. We now discuss some of the possible refinements. In the following, we restrict our attention to the case of bounded rewards (except in Section 2.4.7).

2.4.1 Improved Constants

The regret bound proof for UCB can be improved in two ways. First, the union bound over the different time steps can be replaced by a “peeling” argument. This allows to show a logarithmic regret for any $\alpha > 1$, whereas the proof of Section 2.2 requires $\alpha > 2$ — see [43, Section 2.2] for more details. A second improvement, proposed by Garivier and Cappé [83], is to use a more subtle set of conditions than Equations (2.5)–(2.7). In fact, the authors take both improvements into account, and show that α -UCB has a regret of order $\frac{\alpha}{2} \ln n$ for any $\alpha > 1$. In the limit when α tends to 1, this constant is unimprovable in light of Theorem 2.2 and Equation (2.8).

2.4.2 Second-Order Bounds

Although α -UCB is essentially optimal, the gap between Equation (2.4) and Theorem 2.2 can be important if $\text{kl}(\mu_{i^*}, \mu_i)$ is significantly larger than Δ_i^2 . Several improvements have been proposed toward closing this gap. In particular, the UCB-V algorithm of Audibert et al. [21] takes into account the variance of the distributions and replaces Hoeffding’s inequality by Bernstein’s inequality in the derivation of UCB. A previous algorithm with similar ideas was developed by Auer et al. [23] without theoretical guarantees. A second type of approach replaces L_2 -neighborhoods in α -UCB by kl -neighborhoods. This line of work started with Honda and Takemura [103] where only asymptotic guarantees were

provided. Later, Garivier and Cappé [83] and Maillard et al. [132] (see also Cappé et al., [58]) independently proposed a similar algorithm, called KL-UCB, which is shown to attain the optimal rate in finite-time. More precisely, Garivier and Cappé [83] showed that KL-UCB attains a regret smaller than

$$\sum_{i:\Delta_i>0} \frac{\Delta_i}{\text{kl}(\mu_i, \mu^*)} \alpha \ln n + \mathcal{O}(1)$$

where $\alpha > 1$ is a parameter of the algorithm. Thus, KL-UCB is optimal for Bernoulli distributions and strictly dominates α -UCB for any bounded reward distributions.

2.4.3 Distribution-Free Bounds

In the limit when Δ_i tends to 0, the upper bound in Equation (2.4) becomes vacuous. On the other hand, it is clear that the regret incurred from pulling arm i cannot be larger than $n\Delta_i$. Using this idea, it is easy to show that the regret of α -UCB is always smaller than $\sqrt{\alpha n K \ln n}$ (up to a numerical constant). However, as we shall see in the next section, one can show a minimax lower bound on the regret of order \sqrt{nK} . Audibert and Bubeck [17] proposed a modification of α -UCB that gets rid of the extraneous logarithmic term in the upper bound. More precisely, let $\Delta = \min_{i \neq i^*} \Delta_i$, Audibert and Bubeck [18] show that MOSS (Minimax Optimal Strategy in the Stochastic case) attains a regret smaller than

$$\min \left\{ \sqrt{nK}, \frac{K}{\Delta} \ln \frac{n\Delta^2}{K} \right\}$$

up to a numerical constant. The weakness of this result is that the second term in the above equation only depends on the smallest gap Δ . In Auer and Ortner [25] (see also Perchet and Rigollet, [143]) the authors designed a strategy, called improved UCB, with a regret of order

$$\sum_{i:\Delta_i>0} \frac{1}{\Delta_i} \ln(n\Delta_i^2).$$

This latter regret bound can be better than the one for MOSS in some regimes, but it does not attain the minimax optimal rate of order \sqrt{nK} .

It is an open problem to obtain a strategy with a regret always better than those of MOSS and improved UCB. A plausible conjecture is that a regret of order

$$\sum_{i:\Delta_i>0} \frac{1}{\Delta_i} \ln \frac{n}{H} \quad \text{with } H = \sum_{i:\Delta_i>0} \frac{1}{\Delta_i^2}$$

is attainable. Note that the quantity H appears in other variants of the stochastic multi-armed bandit problem, see Audibert et al. [20].

2.4.4 High Probability Bounds

While bounds on the pseudo-regret \bar{R}_n are important, one typically wants to control the quantity $\hat{R}_n = n\mu^* - \sum_{t=1}^n \mu_{I_t}$ with high probability. Showing that \hat{R}_n is close to its expectation \bar{R}_n is a challenging task, since naively one might expect the fluctuations of \hat{R}_n to be of order \sqrt{n} , which would dominate the expectation \bar{R}_n which is only of order $\ln n$. The concentration properties of \hat{R}_n for UCB are analyzed in detail in Audibert et al. [21], where it is shown that \hat{R}_n concentrates around its expectation, but that there is also a polynomial (in n) probability that \hat{R}_n is of order n . In fact the polynomial concentration of \hat{R}_n around \bar{R}_n can be directly derived from our proof of Theorem 2.1. In Salomon and Audibert [149] it is proved that for anytime strategies (i.e., strategies that do not use the time horizon n) it is basically impossible to improve this polynomial concentration to a classical exponential concentration. In particular this shows that, as far as high probability bounds are concerned, anytime strategies are surprisingly weaker than strategies using the time horizon information (for which exponential concentration of \hat{R}_n around $\ln n$ are possible, see Audibert et al., [21]).

2.4.5 ε -Greedy

A simple and popular heuristic for bandit problems is the ε -greedy strategy — see, e.g., [157]. The idea is very simple. First, pick a parameter $0 < \varepsilon < 1$. Then, at each step greedily play the arm with highest empirical mean reward with probability $1 - \varepsilon$ and play a random arm with probability ε . Auer et al. [23] proved that, if ε is allowed to be a certain function ε_t of the current time step t , namely

$\varepsilon_t = K/(d^2t)$, then the regret grows logarithmically like $(K \ln n)/d^2$, provided $0 < d < \min_{i \neq i^*} \Delta_i$. While this bound has a suboptimal dependence on d , Auer et al. [23] show that this algorithm performs well in practice, but the performance degrades quickly if d is not chosen as a tight lower bound of $\min_{i \neq i^*} \Delta_i$.

2.4.6 Thompson Sampling

In the very first paper on the multi-armed bandit problem, Thompson [162], a simple strategy was proposed for the case of Bernoulli distributions. The so-called Thompson sampling algorithm proceeds as follows. Assume a uniform prior on the parameters $\mu_i \in [0, 1]$, let $\pi_{i,t}$ be the posterior distribution for μ_i at the t -th round, and let $\theta_{i,t} \sim \pi_{i,t}$ (independently from the past given $\pi_{i,t}$). The strategy is then given by $I_t \in \operatorname{argmax}_{i=1, \dots, K} \theta_{i,t}$. Recently there has been a surge of interest for this simple policy, mainly because of its flexibility to incorporate prior knowledge on the arms, see for example Chapelle and Li [65] and the references therein. While the theoretical behavior of Thompson sampling has remained elusive for a long time, we have now a fairly good understanding of its theoretical properties: in Agrawal and Goyal [10] the first logarithmic regret bound was proved, and in Kaufmann et al. [115] it was showed that in fact Thompson sampling attains essentially the same regret than in Equation (2.4). Interestingly note that while Thompson sampling comes from a Bayesian reasoning, it is analyzed with a frequentist perspective. For more on the interplay between Bayesian strategy and frequentist regret analysis we refer the reader to Kaufmann et al. [114].

2.4.7 Heavy-Tailed Distributions

We showed in Section 2.2 how to obtain a UCB-type strategy through a bound on the moment generating function. Moreover one can see that the resulting bound in Theorem 2.1 deteriorates as the tail of the distributions become heavier. In particular, we did not provide any result for the case of distributions for which the moment generating function is not finite. Surprisingly, it was shown in Bubeck et al. [46] that in fact there exists a strategy with essentially the same regret

than in Equation (2.4), as soon as the *variance* of the distributions are finite. More precisely, using more refined robust estimators of the mean than the basic empirical mean, one can construct a UCB-type strategy such that for distributions with moment of order $1 + \varepsilon$ bounded by 1 it satisfies

$$R_n \leq \sum_{i: \Delta_i > 0} \left(8 \left(\frac{4}{\Delta_i} \right)^{\frac{1}{\varepsilon}} \ln n + 5\Delta_i \right).$$

We refer the interested reader to Bubeck et al. [46] for more details on these “robust” versions of UCB.

3

Adversarial Bandits: Fundamental Results

In this section we consider the important variant of the multi-armed bandit problem where no stochastic assumption is made on the generation of rewards. Denote by $g_{i,t}$ the reward (or gain) of arm i at time step t . We assume all rewards are bounded, say $g_{i,t} \in [0, 1]$. At each time step $t = 1, 2, \dots$, simultaneously with the player's choice of the arm $I_t \in \{1, \dots, K\}$, an adversary assigns to each arm $i = 1, \dots, K$ the reward $g_{i,t}$. Similarly to the stochastic setting, we measure the performance of the player compared with the performance of the best arm through the regret

$$R_n = \max_{i=1, \dots, K} \sum_{t=1}^n g_{i,t} - \sum_{t=1}^n g_{I_t,t}.$$

Sometimes we consider losses rather than gains. In this case we denote by $\ell_{i,t}$ the loss of arm i at time step t , and the regret rewrites as

$$R_n = \sum_{t=1}^n \ell_{I_t,t} - \min_{i=1, \dots, K} \sum_{t=1}^n \ell_{i,t}.$$

The loss and gain versions are symmetric, in the sense that one can translate the analysis from one to the other setting via the equivalence

$\ell_{i,t} = 1 - g_{i,t}$. In the following we emphasize the loss version, but we revert to the gain version whenever it makes proofs simpler.

The main goal is to achieve sublinear (in the number of rounds) bounds on the regret uniformly over all possible adversarial assignments of gains to arms. At first sight, this goal might seem hopeless. Indeed, for any deterministic forecaster there exists a sequence of losses $(\ell_{i,t})$ such that $R_n \geq n/2$. Concretely, it suffices to consider the following sequence of losses:

$$\begin{aligned} \text{if } I_t = 1, & \quad \text{then } \ell_{2,t} = 0 \quad \text{and } \ell_{i,t} = 1 \quad \text{for all } i \neq 2; \\ \text{if } I_t \neq 1, & \quad \text{then } \ell_{1,t} = 0 \quad \text{and } \ell_{i,t} = 1 \quad \text{for all } i \neq 1. \end{aligned}$$

The key idea to get around this difficulty is to add randomization to the selection of the action I_t to play. By doing so, the forecaster can “surprise” the adversary, and this surprise effect suffices to get a regret essentially as low as the minimax regret for the stochastic model. Since the regret R_n becomes a random variable, the goal is thus to obtain bounds in high probability or in expectation on R_n (with respect to both eventual randomization of the forecaster and of the adversary). This task is fairly difficult, and a convenient first step is to bound the pseudo-regret

$$\bar{R}_n = \mathbb{E} \sum_{t=1}^n \ell_{I_t,t} - \min_{i=1,\dots,K} \mathbb{E} \sum_{t=1}^n \ell_{i,t}. \quad (3.1)$$

Clearly $\bar{R}_n \leq \mathbb{E} R_n$, and thus an upper bound on the pseudo-regret does not imply a bound on the expected regret. As argued in the Introduction, the pseudo-regret has no natural interpretation unless the adversary is oblivious. In that case, the pseudo-regret coincides with the standard regret, which is always the ultimate quantity of interest.

3.1 Pseudo-regret Bounds

As we pointed out, in order to obtain nontrivial regret guarantees in the adversarial framework it is necessary to consider randomized forecasters. Below we describe the randomized forecaster Exp3, which is based on two fundamental ideas.

Exp3 (Exponential weights for Exploration and Exploitation)

Parameter: a nonincreasing sequence of real numbers $(\eta_t)_{t \in \mathbb{N}}$.

Let p_1 be the uniform distribution over $\{1, \dots, K\}$.

For each round $t = 1, 2, \dots, n$

- (1) Draw an arm I_t from the probability distribution p_t .
- (2) For each arm $i = 1, \dots, K$ compute the estimated loss $\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{p_{i,t}} \mathbb{1}_{I_t=i}$ and update the estimated cumulative loss $\tilde{L}_{i,t} = \tilde{L}_{i,t-1} + \tilde{\ell}_{i,t}$.
- (3) Compute the new probability distribution over arms $p_{t+1} = (p_{1,t+1}, \dots, p_{K,t+1})$, where

$$p_{i,t+1} = \frac{\exp(-\eta_t \tilde{L}_{i,t})}{\sum_{k=1}^K \exp(-\eta_t \tilde{L}_{k,t})}.$$

First, despite the fact that only the loss of the played arm is observed, with a simple trick it is still possible to build an unbiased estimator for the loss of any other arm. Namely, if the next arm I_t to be played is drawn from a probability distribution $p_t = (p_{1,t}, \dots, p_{K,t})$, then

$$\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{p_{i,t}} \mathbb{1}_{I_t=i}$$

is an unbiased estimator (with respect to the draw of I_t) of $\ell_{i,t}$. Indeed, for each $i = 1, \dots, K$ we have

$$\mathbb{E}_{I_t \sim p_t}[\tilde{\ell}_{i,t}] = \sum_{j=1}^K p_{j,t} \frac{\ell_{i,t}}{p_{i,t}} \mathbb{1}_{j=i} = \ell_{i,t}.$$

The second idea is to use an exponential reweighting of the cumulative estimated losses to define the probability distribution p_t from which the forecaster will select the arm I_t . Exponential weighting schemes are a standard tool in the study of sequential prediction schemes under adversarial assumptions. The reader is referred to the monograph by Cesa-Bianchi and Lugosi [61] for a general introduction to prediction of individual sequences, and to the recent monograph by Arora et al. [15] focused on computer science applications of exponential weighting.

We provide two different pseudo-regret bounds for this strategy. The bound (Equation (3.3)) is obtained assuming that the forecaster does

not know the number of rounds n . This is the anytime version of the algorithm. The bound (Equation (3.2)), instead, shows that a better constant can be achieved using the knowledge of the time horizon.

Theorem 3.1 (Pseudo-regret of Exp3). If Exp3 is run with $\eta_t = \eta = \sqrt{\frac{2 \ln K}{nK}}$, then

$$\bar{R}_n \leq \sqrt{2nK \ln K}. \quad (3.2)$$

Moreover, if Exp3 is run with $\eta_t = \sqrt{\frac{\ln K}{tK}}$, then

$$\bar{R}_n \leq 2\sqrt{nK \ln K}. \quad (3.3)$$

Proof. We prove that for any nonincreasing sequence $(\eta_t)_{t \in \mathbb{N}}$ Exp3 satisfies

$$\bar{R}_n \leq \frac{K}{2} \sum_{t=1}^n \eta_t + \frac{\ln K}{\eta_n}. \quad (3.4)$$

Equation (3.2) then trivially follows from Equation (3.4). For Equation (3.3) we use (3.4) and $\sum_{t=1}^n \frac{1}{\sqrt{t}} \leq \int_0^n \frac{1}{\sqrt{t}} dt = 2\sqrt{n}$. The proof of Equation (3.4) is divided into five steps.

First step: Useful equalities

The following equalities can be easily verified:

$$\mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t} = \ell_{I_t,t}, \quad \mathbb{E}_{I_t \sim p_t} \tilde{\ell}_{i,t} = \ell_{i,t}, \quad \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t}^2 = \frac{\ell_{I_t,t}^2}{p_{I_t,t}}, \quad \mathbb{E}_{I_t \sim p_t} \frac{1}{p_{I_t,t}} = K. \quad (3.5)$$

In particular, they imply

$$\sum_{t=1}^n \ell_{I_t,t} - \sum_{t=1}^n \ell_{k,t} = \sum_{t=1}^n \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t} - \sum_{t=1}^n \mathbb{E}_{I_t \sim p_t} \tilde{\ell}_{k,t}. \quad (3.6)$$

The key idea of the proof is to rewrite $\mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t}$ as follows

$$\begin{aligned} \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t} &= \frac{1}{\eta_t} \ln \mathbb{E}_{i \sim p_t} \exp(-\eta_t(\tilde{\ell}_{i,t} - \mathbb{E}_{k \sim p_t} \tilde{\ell}_{k,t})) \\ &\quad - \frac{1}{\eta_t} \ln \mathbb{E}_{i \sim p_t} \exp(-\eta_t \tilde{\ell}_{i,t}). \end{aligned} \quad (3.7)$$

The reader may recognize $\ln \mathbb{E}_{i \sim p_t} \exp(-\eta_t \tilde{\ell}_{i,t})$ as the cumulant-generating function (or the log of the moment-generating function) of the random variable $\tilde{\ell}_{I_t,t}$. This quantity naturally arises in the analysis of forecasters based on exponential weights. In the next two steps we study the two terms in the right-hand side of Equation (3.7).

Second step: Study of the first term in Equation (3.7)

We use the inequalities $\ln x \leq x - 1$ and $\exp(-x) - 1 + x \leq x^2/2$, for all $x \geq 0$, to obtain:

$$\begin{aligned}
& \ln \mathbb{E}_{i \sim p_t} \exp(-\eta_t (\tilde{\ell}_{i,t} - \mathbb{E}_{k \sim p_t} \tilde{\ell}_{k,t})) \\
&= \ln \mathbb{E}_{i \sim p_t} \exp(-\eta_t \tilde{\ell}_{i,t}) + \eta_t \mathbb{E}_{k \sim p_t} \tilde{\ell}_{k,t} \\
&\leq \mathbb{E}_{i \sim p_t} (\exp(-\eta_t \tilde{\ell}_{i,t}) - 1 + \eta_t \tilde{\ell}_{i,t}) \\
&\leq \mathbb{E}_{i \sim p_t} \frac{\eta_t^2 \tilde{\ell}_{i,t}^2}{2} \\
&\leq \frac{\eta_t^2}{2p_{I_t,t}}
\end{aligned} \tag{3.8}$$

where the last step comes from the third equality in Equation (3.5).

Third step: Study of the second term in Equation (3.7)

Let $\tilde{L}_{i,0} = 0$, $\Phi_0(\eta) = 0$ and $\Phi_t(\eta) = \frac{1}{\eta} \ln \frac{1}{K} \sum_{i=1}^K \exp(-\eta \tilde{L}_{i,t})$. Then, by definition of p_t we have

$$\begin{aligned}
-\frac{1}{\eta_t} \ln \mathbb{E}_{i \sim p_t} \exp(-\eta_t \tilde{\ell}_{i,t}) &= -\frac{1}{\eta_t} \ln \frac{\sum_{i=1}^K \exp(-\eta_t \tilde{L}_{i,t})}{\sum_{i=1}^K \exp(-\eta_t \tilde{L}_{i,t-1})} \\
&= \Phi_{t-1}(\eta_t) - \Phi_t(\eta_t).
\end{aligned} \tag{3.9}$$

Fourth step: Summing

Putting together Equations (3.6), (3.7), (3.8), and (3.9) we obtain

$$\sum_{t=1}^n g_{k,t} - \sum_{t=1}^n g_{I_t,t} \leq \sum_{t=1}^n \frac{\eta_t}{2p_{I_t,t}} + \sum_{t=1}^n \Phi_{t-1}(\eta_t) - \Phi_t(\eta_t) - \sum_{t=1}^n \mathbb{E}_{I_t \sim p_t} \tilde{\ell}_{k,t}.$$

The first term is easy to bound in expectation since by the rule of conditional expectations and the last equality in Equation (3.5) we have

$$\mathbb{E} \sum_{t=1}^n \frac{\eta_t}{2p_{I_t,t}} = \mathbb{E} \sum_{t=1}^n \mathbb{E}_{I_t \sim p_t} \frac{\eta_t}{2p_{I_t,t}} = \frac{K}{2} \sum_{t=1}^n \eta_t.$$

For the second term we start with an Abel transformation,

$$\sum_{t=1}^n (\Phi_{t-1}(\eta_t) - \Phi_t(\eta_t)) = \sum_{t=1}^{n-1} (\Phi_t(\eta_{t+1}) - \Phi_t(\eta_t)) - \Phi_n(\eta_n)$$

since $\Phi_0(\eta_1) = 0$. Note that

$$\begin{aligned} -\Phi_n(\eta_n) &= \frac{\ln K}{\eta_n} - \frac{1}{\eta_n} \ln \left(\sum_{i=1}^K \exp(-\eta_n \tilde{L}_{i,n}) \right) \\ &\leq \frac{\ln K}{\eta_n} - \frac{1}{\eta_n} \ln(\exp(-\eta_n \tilde{L}_{k,n})) \\ &= \frac{\ln K}{\eta_n} + \sum_{t=1}^n \tilde{\ell}_{k,t} \end{aligned}$$

and thus we have

$$\mathbb{E} \left[\sum_{t=1}^n g_{k,t} - \sum_{t=1}^n g_{I_t,t} \right] \leq \frac{K}{2} \sum_{t=1}^n \eta_t + \frac{\ln K}{\eta_n} + \mathbb{E} \sum_{t=1}^{n-1} \Phi_t(\eta_{t+1}) - \Phi_t(\eta_t).$$

To conclude the proof, we show that $\Phi'_t(\eta) \geq 0$. Since $\eta_{t+1} \leq \eta_t$, we then obtain $\Phi_t(\eta_{t+1}) - \Phi_t(\eta_t) \leq 0$. Let

$$p_{i,t}^\eta = \frac{\exp(-\eta \tilde{L}_{i,t})}{\sum_{k=1}^K \exp(-\eta \tilde{L}_{k,t})}.$$

Then

$$\begin{aligned} \Phi'_t(\eta) &= -\frac{1}{\eta^2} \ln \left(\frac{1}{K} \sum_{i=1}^K \exp(-\eta \tilde{L}_{i,t}) \right) - \frac{1}{\eta} \frac{\sum_{i=1}^K \tilde{L}_{i,t} \exp(-\eta \tilde{L}_{i,t})}{\sum_{i=1}^K \exp(-\eta \tilde{L}_{i,t})} \\ &= \frac{1}{\eta^2} \frac{1}{\sum_{i=1}^K \exp(-\eta \tilde{L}_{i,t})} \sum_{i=1}^K \exp(-\eta \tilde{L}_{i,t}) \\ &\quad \times \left(-\eta \tilde{L}_{i,t} - \ln \left(\frac{1}{K} \sum_{i=1}^K \exp(-\eta \tilde{L}_{i,t}) \right) \right). \end{aligned}$$

Simplifying, we get (since p_1 is the uniform distribution over $\{1, \dots, K\}$),

$$\Phi'_t(\eta) = \frac{1}{\eta^2} \sum_{i=1}^K p_{i,t}^\eta \ln(K p_{i,t}^\eta) = \frac{1}{\eta^2} \text{KL}(p_t^\eta, p_1) \geq 0. \quad \square$$

3.2 High Probability and Expected Regret Bounds

In this section we prove a high probability bound on the regret. Unfortunately, the Exp3 strategy defined in the previous section is not adequate for this task. Indeed, the variance of the estimate $\tilde{\ell}_{i,t}$ is of order $1/p_{i,t}$, which can be arbitrarily large. In order to ensure that the probabilities $p_{i,t}$ are bounded from below, the original version of Exp3 mixes the exponential weights with a uniform distribution over the arms. In order to avoid increasing the regret, the mixing coefficient γ associated with the uniform distribution cannot be larger than $n^{-1/2}$. Since this implies that the variance of the cumulative loss estimate $\tilde{L}_{i,n}$ can be of order $n^{3/2}$, very little can be said about the concentration of the regret also for this variant of Exp3.

This issue can be solved by combining the mixing idea with a different estimate for losses. In fact, the core idea is more transparent when expressed in terms of gains, and so we turn to the gain version of the problem. The trick is to introduce a bias in the gain estimate which allows to derive a high probability statement on this estimate.

Lemma 3.2. For $\beta \leq 1$, let

$$\tilde{g}_{i,t} = \frac{g_{i,t} \mathbf{1}_{I_t=i} + \beta}{p_{i,t}}.$$

Then, with probability at least $1 - \delta$,

$$\sum_{t=1}^n g_{i,t} \leq \sum_{t=1}^n \tilde{g}_{i,t} + \frac{\ln(\delta^{-1})}{\beta}.$$

Proof. Let \mathbb{E}_t be the expectation conditioned on I_1, \dots, I_{t-1} . Since $\exp(x) \leq 1 + x + x^2$ for $x \leq 1$, for $\beta \leq 1$ we have

$$\begin{aligned}
& \mathbb{E}_t \exp \left(\beta g_{i,t} - \beta \frac{g_{i,t} \mathbb{1}_{I_t=i} + \beta}{p_{i,t}} \right) \\
& \leq \left(1 + \mathbb{E}_t \left[\beta g_{i,t} - \beta \frac{g_{i,t} \mathbb{1}_{I_t=i}}{p_{i,t}} \right] + \mathbb{E}_t \left[\beta g_{i,t} - \beta \frac{g_{i,t} \mathbb{1}_{I_t=i}}{p_{i,t}} \right]^2 \right) \\
& \quad \times \exp \left(-\frac{\beta^2}{p_{i,t}} \right) \\
& \leq \left(1 + \beta^2 \frac{g_{i,t}^2}{p_{i,t}} \right) \exp \left(-\frac{\beta^2}{p_{i,t}} \right) \\
& \leq 1
\end{aligned}$$

where the last inequality uses $1 + u \leq \exp(u)$. As a consequence, we have

$$\mathbb{E} \exp \left(\beta \sum_{t=1}^n g_{i,t} - \beta \sum_{t=1}^n \frac{g_{i,t} \mathbb{1}_{I_t=i} + \beta}{p_{i,t}} \right) \leq 1.$$

Moreover, Markov's inequality implies $\mathbb{P}(X > \ln(\delta^{-1})) \leq \delta \mathbb{E}e^X$ and thus, with probability at least $1 - \delta$,

$$\beta \sum_{t=1}^n g_{i,t} - \beta \sum_{t=1}^n \frac{g_{i,t} \mathbb{1}_{I_t=i} + \beta}{p_{i,t}} \leq \ln(\delta^{-1}). \quad \square$$

The strategy associated with these new estimates, called Exp3.P, is described in Figure 3.1. Note that, for the sake of simplicity, the strategy is described in the setting with known time horizon (η is constant). Anytime results can easily be derived with the same techniques as in the proof of Theorem 3.1.

In the next theorem we propose two different high probability bounds. In Equation (3.10) the algorithm needs the confidence level δ as an input parameter. In Equation (3.11) the algorithm satisfies a high probability bound for any confidence level. This latter property is particularly important to derive good bounds on the expected regret.

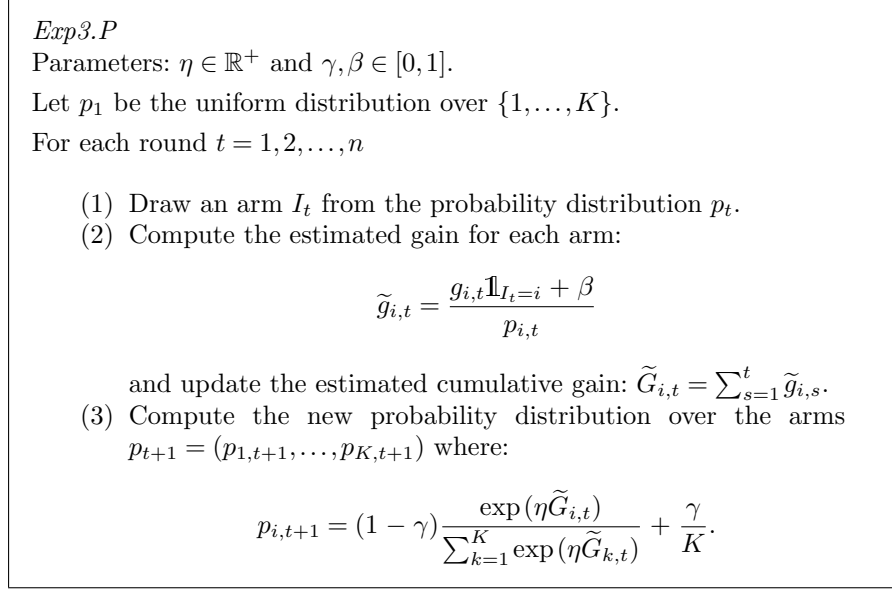


Fig. 3.1 Exp3.P forecaster.

Theorem 3.3 (High probability bound for Exp3.P). For any given $\delta \in (0, 1)$, if Exp3.P is run with

$$\beta = \sqrt{\frac{\ln(K\delta^{-1})}{nK}}, \quad \eta = 0.95 \sqrt{\frac{\ln(K)}{nK}}, \quad \gamma = 1.05 \sqrt{\frac{K \ln(K)}{n}}$$

then, with probability at least $1 - \delta$,

$$R_n \leq 5.15 \sqrt{nK \ln(K\delta^{-1})}. \quad (3.10)$$

Moreover, if Exp3.P is run with $\beta = \sqrt{\frac{\ln(K)}{nK}}$, while η and γ are chosen as before, then, with probability at least $1 - \delta$,

$$R_n \leq \sqrt{\frac{nK}{\ln(K)} \ln(\delta^{-1})} + 5.15 \sqrt{nK \ln(K)}. \quad (3.11)$$

Proof. We first prove (in three steps) that if $\gamma \leq 1/2$ and $(1 + \beta)K\eta \leq \gamma$, then Exp3.P satisfies, with probability at least $1 - \delta$,

$$R_n \leq \beta nK + \gamma n + (1 + \beta)\eta K n + \frac{\ln(K\delta^{-1})}{\beta} + \frac{\ln K}{\eta}. \quad (3.12)$$

First step: Notation and simple equalities

One can immediately see that $\mathbb{E}_{i \sim p_t} \tilde{g}_{i,t} = g_{I_t,t} + \beta K$, and thus

$$\sum_{t=1}^n g_{k,t} - \sum_{t=1}^n g_{I_t,t} = \beta n K + \sum_{t=1}^n g_{k,t} - \sum_{t=1}^n \mathbb{E}_{i \sim p_t} \tilde{g}_{i,t}. \quad (3.13)$$

The key step is, again, to consider the cumulant-generating function of $\tilde{g}_{i,t}$. However, because of the mixing, we need to introduce a few more notations. Let $u = (\frac{1}{K}, \dots, \frac{1}{K})$ be the uniform distribution over the arms, let and $w_t = \frac{p_t - u}{1-\gamma}$ be the distribution induced by Exp3.P at time t without the mixing. Then we have:

$$\begin{aligned} -\mathbb{E}_{i \sim p_t} \tilde{g}_{i,t} &= -(1-\gamma) \mathbb{E}_{i \sim w_t} \tilde{g}_{i,t} - \gamma \mathbb{E}_{i \sim u} \tilde{g}_{i,t} \\ &= (1-\gamma) \left(\frac{1}{\eta} \ln \mathbb{E}_{i \sim w_t} \exp(\eta(\tilde{g}_{i,t} - \mathbb{E}_{k \sim w_t} \tilde{g}_{k,t})) \right. \\ &\quad \left. - \frac{1}{\eta} \ln \mathbb{E}_{i \sim w_t} \exp(\eta \tilde{g}_{i,t}) \right) - \gamma \mathbb{E}_{i \sim u} \tilde{g}_{i,t}. \end{aligned} \quad (3.14)$$

Second step: Study of the first term in Equation (3.14)

We use the inequalities $\ln x \leq x - 1$ and $\exp(x) \leq 1 + x + x^2$, for all $x \leq 1$, as well as the fact that $\eta \tilde{g}_{i,t} \leq 1$ since $(1 + \beta)\eta K \leq \gamma$:

$$\begin{aligned} \ln \mathbb{E}_{i \sim w_t} \exp(\eta(\tilde{g}_{i,t} - \mathbb{E}_{k \sim p_t} \tilde{g}_{k,t})) &= \ln \mathbb{E}_{i \sim w_t} \exp(\eta \tilde{g}_{i,t}) - \eta \mathbb{E}_{k \sim p_t} \tilde{g}_{k,t} \\ &\leq \mathbb{E}_{i \sim w_t} [\exp(\eta \tilde{g}_{i,t}) - 1 - \eta \tilde{g}_{i,t}] \\ &\leq \mathbb{E}_{i \sim w_t} \eta^2 \tilde{g}_{i,t}^2 \\ &\leq \frac{1 + \beta}{1 - \gamma} \eta^2 \sum_{i=1}^K \tilde{g}_{i,t} \end{aligned} \quad (3.15)$$

where we used $\frac{w_{i,t}}{p_{i,t}} \leq \frac{1}{1-\gamma}$ in the last step.

Third step: Summing

Set $\tilde{G}_{i,0} = 0$. Recall that $w_t = (w_{1,t}, \dots, w_{K,t})$ with

$$w_{i,t} = \frac{\exp(-\eta \tilde{G}_{i,t-1})}{\sum_{k=1}^K \exp(-\eta \tilde{G}_{k,t-1})}. \quad (3.16)$$

Then substituting Equation (3.15) into Equation (3.14) and summing using Equation (3.16), we obtain

$$\begin{aligned}
& - \sum_{t=1}^n \mathbb{E}_{i \sim p_t} \tilde{g}_{i,t} \\
& \leq (1 + \beta)\eta \sum_{t=1}^n \sum_{i=1}^K \tilde{g}_{i,t} - \frac{1 - \gamma}{\eta} \sum_{t=1}^n \ln \left(\sum_{i=1}^K w_{i,t} \exp(\eta \tilde{g}_{i,t}) \right) \\
& = (1 + \beta)\eta \sum_{t=1}^n \sum_{i=1}^K \tilde{g}_{i,t} - \frac{1 - \gamma}{\eta} \ln \left(\prod_{t=1}^n \frac{\sum_{i=1}^K \exp(\eta \tilde{G}_{i,t})}{\sum_{i=1}^K \exp(\eta \tilde{G}_{i,t-1})} \right) \\
& \leq (1 + \beta)\eta K \max_j \tilde{G}_{j,n} + \frac{\ln K}{\eta} - \frac{1 - \gamma}{\eta} \ln \left(\sum_{t=1}^n \exp(\eta \tilde{G}_{i,n}) \right) \\
& \leq -(1 - \gamma - (1 + \beta)\eta K) \max_j \tilde{G}_{j,n} + \frac{\ln(K)}{\eta} \\
& \leq -(1 - \gamma - (1 + \beta)\eta K) \max_j \sum_{t=1}^n g_{j,t} + \frac{\ln(K\delta^{-1})}{\beta} + \frac{\ln(K)}{\eta}.
\end{aligned}$$

The last inequality comes from Lemma 3.2, the union bound, and $\gamma - (1 + \beta)\eta K \leq 1$ which is a consequence of $(1 + \beta)\eta K \leq \gamma \leq 1/2$. Combining this last inequality with Equation (3.13) we obtain

$$R_n \leq \beta n K + \gamma n + (1 + \beta)\eta K n + \frac{\ln(K\delta^{-1})}{\beta} + \frac{\ln(K)}{\eta}$$

which is the desired result.

Equation (3.10) is then proved as follows. First, it is trivial if $n \geq 5.15\sqrt{nK \ln(K\delta^{-1})}$ and thus we can assume that this is not the case. This implies that $\gamma \leq 0.21$ and $\beta \leq 0.1$, and thus we have $(1 + \beta)\eta K \leq \gamma \leq 1/2$. Using Equation (3.12) directly yields the claimed bound. The same argument can be used to derive Equation (3.11). \square

We now discuss the expected regret bounds. As the cautious reader may already have observed, if the adversary is oblivious, namely when $(\ell_{1,t}, \dots, \ell_{K,t})$ is independent of I_1, \dots, I_{t-1} for each t , a pseudo-regret bound implies the same bound on the expected regret. This follows from

noting that the expected regret against an oblivious adversary is smaller than the maximal pseudo-regret against deterministic adversaries, see [18, Proposition 33] for a proof of this fact. In the general case of a nonoblivious adversary, the loss vector $(\ell_{1,t}, \dots, \ell_{K,t})$ at time t depends on the past actions of the forecaster. This makes the analysis of the expected regret more intricate. One way around this difficulty is to first prove high probability bounds and then integrate the resulting bound. Following this method, we derive a bound on the expected regret of Exp3.P using Equation (3.11).

Theorem 3.4 (Expected regret of Exp3.P). If Exp3.P is run with

$$\beta = \sqrt{\frac{\ln K}{nK}}, \quad \eta = 0.95\sqrt{\frac{\ln K}{nK}}, \quad \gamma = 1.05\sqrt{\frac{K \ln K}{n}}$$

then

$$\mathbb{E} R_n \leq 5.15\sqrt{nK \ln K} + \sqrt{\frac{nK}{\ln K}}. \quad (3.17)$$

Proof. We integrate the deviations in Equation (3.11) using the formula

$$\mathbb{E} W \leq \int_0^1 \frac{1}{\delta} \mathbb{P} \left(W > \ln \frac{1}{\delta} \right) d\delta$$

for a real-valued random variable W . In particular, taking

$$W = \sqrt{\frac{\ln K}{nK}} (R_n - 5.15\sqrt{nK \ln K})$$

yields $\mathbb{E} W \leq 1$, which is equivalent to Equation (3.17). \square

3.3 Lower Bound

The next theorem shows that the results of the previous sections are essentially unimprovable, up to logarithmic factors. The result is proven via the probabilistic method: we show that there exists a distribution of rewards for the arms such that the pseudo-regret of any forecaster must be high when averaged over this distribution. Owing to this probabilistic construction, the lower bound proof is based on the same Kullback–Leibler divergence as the one used in the proof of the lower bound

for stochastic bandits — see Section 2.3. We are not aware of other techniques for proving bandit lower bounds.

We find it more convenient to prove the results for rewards rather than losses. In order to emphasize that our rewards are stochastic (in particular, Bernoulli random variables), we use $Y_{i,t} \in \{0,1\}$ to denote the reward obtained by pulling arm i at time t .

Theorem 3.5 (Minimax lower bound). Let \sup be the supremum over all distribution of rewards such that, for $i = 1, \dots, K$, the rewards $Y_{i,1}, Y_{i,2}, \dots \in \{0,1\}$ are i.i.d., and let \inf be the infimum over all forecasters. Then

$$\inf \sup \left(\max_{i=1, \dots, K} \mathbb{E} \sum_{t=1}^n Y_{i,t} - \mathbb{E} \sum_{t=1}^n Y_{I_t,t} \right) \geq \frac{1}{20} \sqrt{nK} \quad (3.18)$$

where expectations are with respect to both the random generation of rewards and the internal randomization of the forecaster.

Since $\max_{i=1, \dots, K} \mathbb{E} \sum_{t=1}^n Y_{i,t} - \mathbb{E} \sum_{t=1}^n Y_{I_t,t} = \bar{R}_n \leq \mathbb{E} R_n$, Theorem 3.5 immediately entails a lower bound on the regret of any forecaster.

The general idea of the proof goes as follows. Since at least one arm is pulled less than n/K times, for this arm one cannot differentiate between a Bernoulli of parameter $1/2$ and a Bernoulli of parameter $1/2 + \sqrt{K/n}$. Thus, if all arms are Bernoulli of parameter $1/2$, but one whose parameter is $1/2 + \sqrt{K/n}$, then the forecaster should incur a regret of order $n\sqrt{K/n} = \sqrt{nK}$. In order to formalize this idea, we use the Kullback–Leibler divergence, and in particular Pinsker’s inequality, to compare the behavior of a given forecaster against: (1) the distribution where all arms are Bernoulli of parameter $1/2$, and (2) the same distribution where the parameter of one arm is increased by ε .

We start by proving a more general lemma, which could also be used to derive lower bounds in other contexts. The proof of Theorem 3.5 then follows by a simple optimization over ε .

Lemma 3.6. Let $\varepsilon \in [0,1)$. For any $i \in \{1, \dots, K\}$ let \mathbb{E}_i be the expectation against the joint distribution of rewards where all arms are i.i.d.

Bernoulli of parameter $\frac{1-\varepsilon}{2}$ but arm i , which is i.i.d. Bernoulli of parameter $\frac{1+\varepsilon}{2}$. Then, for any forecaster,

$$\max_{i=1,\dots,K} \mathbb{E}_i \sum_{t=1}^n (Y_{i,t} - Y_{I_t,t}) \geq n\varepsilon \left(1 - \frac{1}{K} - \sqrt{\varepsilon \ln \frac{1+\varepsilon}{1-\varepsilon}} \sqrt{\frac{n}{2K}} \right).$$

Proof. We provide a proof in five steps by lower bounding $\frac{1}{K} \sum_{i=1}^K \mathbb{E}_i \sum_{t=1}^n (Y_{i,t} - Y_{I_t,t})$. This implies the statement of the lemma because a max is larger than a mean.

First step: Empirical distribution of plays

We start by considering a deterministic forecaster. Let $q_n = (q_{1,n}, \dots, q_{K,n})$ be the empirical distribution of plays over the arms defined by $q_{i,n} = \frac{T_i(n)}{n}$ — recall from Section 2 that $T_i(n)$ denotes the number of times arm i was selected in the first n rounds. Let J_n be drawn according to q_n . We denote by \mathbb{P}_i the law of J_n against the distribution where all arms are i.i.d. Bernoulli of parameter $\frac{1-\varepsilon}{2}$ but arm i , which is i.i.d. Bernoulli of parameter $\frac{1+\varepsilon}{2}$ (we call this the i -th stochastic adversary). Recall that $\mathbb{P}_i(J_n = j) = \mathbb{E}_i \frac{T_j(n)}{n}$, hence

$$\mathbb{E}_i \sum_{t=1}^n (Y_{i,t} - Y_{I_t,t}) = \varepsilon n \sum_{j \neq i} \mathbb{P}_i(J_n = j) = \varepsilon n (1 - \mathbb{P}_i(J_n = i))$$

which implies

$$\frac{1}{K} \sum_{i=1}^K \mathbb{E}_i \sum_{t=1}^n (Y_{i,t} - Y_{I_t,t}) = \varepsilon n \left(1 - \frac{1}{K} \sum_{i=1}^K \mathbb{P}_i(J_n = i) \right). \quad (3.19)$$

Second step: Pinsker's inequality

Let \mathbb{P}_0 be the law of J_n for the distribution where all arms are i.i.d. Bernoulli of parameter $\frac{1-\varepsilon}{2}$. Then Pinsker's inequality immediately gives $\mathbb{P}_i(J_n = i) \leq \mathbb{P}_0(J_n = i) + \sqrt{\frac{1}{2} \text{KL}(\mathbb{P}_0, \mathbb{P}_i)}$, and so

$$\frac{1}{K} \sum_{i=1}^K \mathbb{P}_i(J_n = i) \leq \frac{1}{K} + \frac{1}{K} \sum_{i=1}^K \sqrt{\frac{1}{2} \text{KL}(\mathbb{P}_0, \mathbb{P}_i)}. \quad (3.20)$$

Third step: Computation of $\text{KL}(\mathbb{P}_0, \mathbb{P}_i)$

Since the forecaster is deterministic, the sequence of rewards $Y^n = (Y_1, \dots, Y_n) \in \{0, 1\}^n$ received by the forecaster uniquely determines the empirical distribution of plays q_n . In particular, the law of J_n conditionally to Y^n is the same for any i -th stochastic adversary. For each $i = 0, \dots, K$, let \mathbb{P}_i^n be the law of Y^n against the i -th adversary. Then one can easily show that $\text{KL}(\mathbb{P}_0, \mathbb{P}_i) \leq \text{KL}(\mathbb{P}_0^n, \mathbb{P}_i^n)$. Now we use the chain rule for Kullback–Leibler divergence — see for example [61, Section A.2] — iteratively to introduce the laws \mathbb{P}_i^t of $Y^t = (Y_1, \dots, Y_t)$. More precisely, we have

$$\begin{aligned}
& \text{KL}(\mathbb{P}_0^n, \mathbb{P}_i^n) \\
&= \text{KL}(\mathbb{P}_0^1, \mathbb{P}_i^1) + \sum_{t=2}^n \sum_{y^{t-1}} \mathbb{P}_0^{t-1}(y^{t-1}) \text{KL}(\mathbb{P}_0^t(\cdot | y^{t-1}), \mathbb{P}_i^t(\cdot | y^{t-1})) \\
&= \text{KL}(\mathbb{P}_0^1, \mathbb{P}_i^1) + \sum_{t=2}^n \left(\sum_{y^{t-1}: I_t=i} \mathbb{P}_0^{t-1}(y^{t-1}) \text{KL}\left(\frac{1-\varepsilon}{2}, \frac{1+\varepsilon}{2}\right) \right. \\
&\quad \left. + \sum_{y^{t-1}: I_t \neq i} \mathbb{P}_0^{t-1}(y^{t-1}) \text{KL}\left(\frac{1+\varepsilon}{2}, \frac{1+\varepsilon}{2}\right) \right) \\
&= \text{KL}\left(\frac{1-\varepsilon}{2}, \frac{1+\varepsilon}{2}\right) \mathbb{E}_0 T_i(n). \tag{3.21}
\end{aligned}$$

Fourth step: Conclusion for deterministic forecasters

By using that the square root is concave and combining $\text{KL}(\mathbb{P}_0, \mathbb{P}_i) \leq \text{KL}(\mathbb{P}_0^n, \mathbb{P}_i^n)$ with Equation (3.21), we deduce that

$$\begin{aligned}
\frac{1}{K} \sum_{i=1}^K \sqrt{\text{KL}(\mathbb{P}_0, \mathbb{P}_i)} &\leq \sqrt{\frac{1}{K} \sum_{i=1}^K \text{KL}(\mathbb{P}_0, \mathbb{P}_i)} \\
&\leq \sqrt{\frac{1}{K} \sum_{i=1}^K \text{KL}\left(\frac{1-\varepsilon}{2}, \frac{1+\varepsilon}{2}\right) \mathbb{E}_0 T_i(n)} \\
&= \sqrt{\frac{n}{K} \text{KL}\left(\frac{1-\varepsilon}{2}, \frac{1+\varepsilon}{2}\right)}. \tag{3.22}
\end{aligned}$$

We conclude the proof for deterministic forecasters by applying Equations (3.20) and (3.22) to Equation (3.19), and observing that $\text{KL}(\frac{1-\varepsilon}{2}, \frac{1+\varepsilon}{2}) = \varepsilon \ln \frac{1+\varepsilon}{1-\varepsilon}$.

Fifth step: Randomized forecasters via Fubini's Theorem

Extending previous results to randomized forecasters is easy. Denote by \mathbb{E}_r the expectation with respect to the forecaster's internal randomization. Then Fubini's Theorem implies

$$\frac{1}{K} \sum_{i=1}^K \mathbb{E}_i \sum_{t=1}^n \mathbb{E}_r(Y_{i,t} - Y_{I_t,t}) = \mathbb{E}_r \frac{1}{K} \sum_{i=1}^K \mathbb{E}_i \sum_{t=1}^n (Y_{i,t} - Y_{I_t,t}).$$

Now the proof is concluded by applying the lower bound on $\frac{1}{K} \sum_{i=1}^K \mathbb{E}_i \sum_{t=1}^n (Y_{i,t} - Y_{I_t,t})$, which we proved in previous steps, to each realization of the forecaster's random bits. \square

3.4 Refinements and Bibliographic Remarks

The adversarial framework studied in this section was originally investigated in a *full information* setting, where at the end of each round the forecaster observes the complete loss vector $(\ell_{1,t}, \dots, \ell_{K,t})$. We refer the reader to [61] for the history of this problem. The Exp3 and Exp3.P strategies were introduced¹ and analyzed by Auer et al. [24], where the lower bound of Theorem 3.5 is also proven. The proofs presented in this section are taken from Ref. [43]. We now give an overview of some of the many improvements and refinements that have been proposed since these initial analyses.

3.4.1 Log-Free Upper Bounds

One can see that there is a logarithmic gap between the pseudo-regret of Exp3, presented in Theorem 3.1, and the minimax lower bound of Theorem 3.5. This gap was closed by Audibert and Bubeck [17], who constructed a new class of strategies and showed that some of them

¹In its original formulation the Exp3 strategy was defined as a mixture of exponential weights with the uniform distribution on the set of arms. It was noted in Stoltz [156] that this mixing is not necessary, see footnote 2 on page 26 in Bubeck [43] for more details on this.

have a pseudo-regret of order \sqrt{nK} . This new class of strategies, called INF (Implicitly Normalized Forecaster), is based on the following idea. First, note that one can generalize the exponential weighting scheme of Exp3 as follows: given a potential function ψ , assign the probability

$$p_{i,t+1} = \frac{\psi(\tilde{L}_{i,t})}{\sum_{j=1}^K \psi(\tilde{L}_{j,t})}.$$

This type of strategy is called a weighted average forecaster, see [61, Section 2]. In INF the normalization is done implicitly, by a translation of the losses. More precisely, INF with potential ψ assigns the probability $p_{i,t+1} = \psi(C_t - \tilde{L}_{i,t})$, where C_t is the constant such that p_{t+1} sum to 1. The key to obtain a minimax optimal pseudo-regret is to take ψ of the form $\psi(x) = (-\eta x)^{-q}$ with $q > 1$, while Exp3 corresponds to $\psi(x) = \exp(\eta x)$. Audibert et al. [19] realized that the INF strategy can be formulated as a Mirror Descent algorithm. This point of view significantly simplifies the proofs. We refer the reader to Section 5 (and in particular Theorem 5.10) for more details.

While it is possible to get log-free pseudo-regret bounds, the situation becomes significantly more complicated when one considers high probability regret and expected regret. Audibert and Bubeck [18] proved that one can get a log-free expected regret if the adversary is oblivious, i.e., the sequence of loss vectors is independent of the forecaster's actions. Moreover, it is also possible to get a log-free high probability regret if the adversary is fully oblivious (i.e., the loss vectors are independently drawn, but not identically distributed — this includes the oblivious adversary). It is conjectured [18] that it is not possible to obtain a log-free expected regret bound against a general nonoblivious adversary.

3.4.2 Adaptive Bounds

One of the strengths of the bounds proposed in this section is also one of its weaknesses: the bounds hold against any adversary. It is clear that in some cases it is possible to obtain a much smaller regret than the worst case regret. For example, when the sequence of losses is an i.i.d. sequence, we proved in Section 2 that it is possible to obtain a

logarithmic pseudo-regret (provided that the gap Δ is considered as a constant). Thus, it is natural to ask if it is possible to have strategies with minimax optimal regret, but also with much smaller regret when the loss sequence is not worst case.

The first bound in this direction was proved by Auer et al. [24], who showed that, for the gain version of the problem and against an oblivious adversary, Exp3 has a pseudo-regret of order $\sqrt{KG_n^*}$ (omitting log factors), where $G_n^* \leq n$ is the maximal cumulative reward of the optimal arm after n rounds. This result was improved by Audibert and Bubeck [18], who showed that using the gain estimate

$$\tilde{g}_{i,t} = -\frac{\mathbb{1}_{I_t=i}}{\beta} \ln \left(1 - \frac{\beta g_{i,t}}{p_{i,t}} \right)$$

one can bound the regret with high probability by essentially the same quantity as before, and against any adversary.

Another direction was explored by Hazan and Kale [96] building on previous works in the full information setting — [64]. In this work the authors proved that one can attain a regret of order $\sqrt{\sum_{i=1}^K V_{i,n}}$ excluding log factors, where

$$V_{i,n} = \sum_{t=1}^n \left(\ell_{i,t} - \frac{1}{n} \sum_{s=1}^n \ell_{i,s} \right)^2$$

is the total variation of the loss for arm i . In fact their result is more general, as it applies to the linear bandit framework — see Section 5. The main new ingredient in their analysis is a “reservoir sampling” procedure. We refer the reader to [96] for details. See also the works of Slivkins and Upfal [153] and Slivkins [152] for related results on slowly changing bandits.

In Section 3.4.4 below we describe another type of adaptive bound, where one combines minimax optimal regret for the adversarial model with logarithmic pseudo-regret for the stochastic model.

3.4.3 Competing with the Best Switching Strategy

While competing against the policy consistently playing the best fixed arm is a natural way of defining regret, in some applications it might be

interesting to consider regret with respect to a bigger class of policies. Though this problem is the focus of Section 4, there is a class of natural policies that can be directly dealt with by the methods of this section. Namely, consider the problem of competing against any policy constrained to make at most $S \leq n$ switches (a switch is when the arm played at time t is different from the arm played at time $t + 1$). This problem was studied by Auer [22], where it was first shown that a simple variant of Exp3 attains a low switching regret against oblivious adversaries. Later, Audibert and Bubeck [18] proved that Exp3.P attains an expected regret (and a high probability regret) of order $\sqrt{nKS \ln(nK/S)}$ for this problem.

3.4.4 Stochastic versus Adversarial Bandits

From a practical viewpoint, Exp3 should be a safe choice when we have reasons to believe that the sequence of rewards is not well matched by any i.i.d. process. Indeed, it is easy to prove that UCB can have linear regret, i.e., $\bar{R}_n = \Omega(n)$, on certain deterministic sequences. In [53] a new strategy was described, called SAO (Stochastic and Adversarial Optimal), which enjoys (up to logarithmic factors) both the guarantee of Exp3 for the adversarial model and the guarantee of UCB for the stochastic model. More precisely SAO satisfies $\bar{R}_n = \mathcal{O}(\frac{K}{\Delta} \log^2(n) \log(K))$ in the stochastic model and $\bar{R}_n = \mathcal{O}(\sqrt{nK} \log^{3/2}(n) \log(K))$ in the adversarial model. Note that while this result is a step toward more flexible strategies, the very notion of regret \bar{R}_n can become vacuous with nonstationarities in the reward sequence, since the total reward of the best fixed action might be very small. In that case the notion of switching regret — see Section 3.4.3 — is more relevant, and it would be interesting to derive a strategy with logarithmic regret in the stochastic model, and a switching regret of order \sqrt{nKS} in the adversarial model.

3.4.5 Alternative Feedback Structures

As mentioned at the beginning of this section, the adversarial multi-armed bandit is a variation of the full information setting, with a weaker feedback signal (only the incurred loss versus the full vector of losses is

observed). Many other feedback structures can be considered, and we conclude the section by describing a few of them.

In the *label efficient* setting, originally proposed by Helmbold and Panizza [99], at the end of each round the forecaster has to decide whether to ask for the losses of the current round, knowing that this can be done for at most $m \leq n$ times. In this setting, Cesa-Bianchi et al. [63] proved that the minimax pseudo-regret is of order $n\sqrt{\frac{\ln K}{m}}$. A bandit label efficient version was proposed by Allenberg et al. [11]. Audibert and Bubeck [18] proved that the minimax pseudo-regret for the bandit label efficient version is of order $n\sqrt{\frac{K}{m}}$. These results do not require any fundamentally new algorithmic idea, besides the fact the forecaster has to randomize to select the rounds in which the losses are revealed. Roughly speaking, a simple coin toss with parameter $\varepsilon = m/n$ is sufficient to obtain an optimal regret.

Mannor and Shamir [133] study a model that interpolates between the full information and the bandit setting. The basic idea is that there is an undirected graph G with K vertices (one vertex for each arm) that encodes the feedback structure. When one pulls arm i the losses of all neighboring arms $j \in N(i)$ in the graph are observed. Thus, a graph with no edges is equivalent to the bandit problem, while the complete graph is equivalent to the full information setting. Given the feedback structure G , it is natural to consider the following unbiased loss estimate

$$\tilde{\ell}_{i,t} = \frac{\ell_{i,t} \mathbb{1}_{i \in N(I_t)}}{\sum_{j \in N(i)} p_{j,t}}.$$

Using Exp3 with this loss estimate, the authors show that the minimax pseudo-regret (up to logarithmic factors) is of order of $\sqrt{\alpha(G)n}$, where $\alpha(G)$ is the independence number of graph G . Note that this interpolated setting naturally arises in applications like ad placement on websites. Indeed, if a user clicks on an advertisement, it is plausible to assume that the same user would have clicked on similar advertisements, had they been displayed.

The above problems are all specific examples of the more general partial monitoring setting. In this model, at the end of each round the

player does not observe the incurred loss $\ell_{I_t,t}$ but rather a stochastic “signal” $S_{I_t,t}$. A prototypical example of this scenario is the following: a website is repeatedly selling the same item to a sequence of customers. The selling price is dynamically adjusted, and each customer buys the item only if the current price is smaller or equal than his own hidden value for the item. The pricing algorithm (i.e., the player in our terminology) does not see each user’s value, but only whether the user bought the item or not.

The relationship between the signals and the incurred losses defines the instance of a partial monitoring problem. We refer the interested reader to [61] for more details, including an historical account. Recent progress on this problem has been made by Bartok et al. [34, 35] and Foster and Rakhlin [80].

4

Contextual Bandits

A natural extension of the multi-armed problem is obtained by associating side information with each arm. Based on this side information, or context, a notion of “contextual regret” is introduced where optimality is defined with respect to the best policy (i.e., mapping from contexts to arms) rather than the best arm. The space of policies, within which the optimum is sought, is typically chosen in order to have some desired structure. A different viewpoint is obtained when the contexts are privately accessed by the policies (which are then appropriately called “experts”). In this case the contextual information is hidden from the forecaster, and arms must be chosen based only on the past estimated performance of the experts.

Contextual bandits naturally arise in many applications. For example, in personalized news article recommendation the task is to select, from a pool of candidates, a news article to display whenever a new user visits a website. The articles correspond to arms, and a reward is obtained whenever the user clicks on the selected article. Side information, in the form of features, can be extracted from both user and articles. For the user this may include historical activities, demographic information, and geolocation; for the articles, we may have content

information and categories. See [128] for more details on this application of contextual bandits.

In general, the presence of contexts creates a wide spectrum of possible variations obtained by combining assumptions on the rewards with assumptions on the nature of contexts and policies. In this section we describe just a few of the results available in the literature and use the bibliographic remarks to mention all those that we are aware of.

4.1 Bandits with Side Information

The most basic example of contextual bandits is obtained when game rounds $t = 1, 2, \dots$ are marked by contexts s_1, s_2, \dots from a given context set \mathcal{S} . The forecaster must learn the best mapping $g : \mathcal{S} \rightarrow \{1, \dots, K\}$ of contexts to arms. We analyze this simple side information setting in the case of adversarial rewards, and we further assume that the sequence of contexts s_t is arbitrary but fixed. The approach we take is the simplest: run a separate instance of Exp3 on each distinct context.

We introduce the following notion of pseudo-regret

$$\bar{R}_n^{\mathcal{S}} = \max_{g: \mathcal{S} \rightarrow \{1, \dots, K\}} \mathbb{E} \left[\sum_{t=1}^n \ell_{I_t, t} - \sum_{t=1}^n \ell_{g(s_t), t} \right].$$

Here $s_t \in \mathcal{S}$ denotes the context marking the t -th game round. A bound on this pseudoregret is almost immediately obtained using the adversarial bandit results from Section 3.

Theorem 4.1. There exists a randomized forecaster for bandits with side information (the \mathcal{S} -Exp3 forecaster, defined in the proof) that satisfies

$$\bar{R}_n^{\mathcal{S}} \leq \sqrt{2n|\mathcal{S}|K \ln K}$$

for any set \mathcal{S} of contexts.

Proof. Let $S = |\mathcal{S}|$. The \mathcal{S} -Exp3 forecaster runs an instance of Exp3 on each context $s \in \mathcal{S}$. Let n_s the number of times when $s_t = s$ within the first n time steps. Using the bound (Equation (3.2)) established in

Theorem 3.1 we get

$$\begin{aligned} \max_{g: \mathcal{S} \rightarrow \{1, \dots, K\}} \mathbb{E} \left[\sum_{t=1}^n (\ell_{I_t, t} - \ell_{g(s_t), t}) \right] &= \sum_{s \in \mathcal{S}} \max_{k=1, \dots, K} \mathbb{E} \left[\sum_{t: s_t=s} (\ell_{I_t, t} - \ell_{k, t}) \right] \\ &\leq \sum_{s \in \mathcal{S}} \sqrt{2n_s K \ln K} \\ &\leq \sqrt{2nSK \ln K} \end{aligned}$$

where in the last step we used Jensen's inequality and the identity $\sum_s n_s = n$. \square

In Section 4.2.1, we extend this construction by considering several context sets simultaneously.

A lower bound $\Omega(\sqrt{nSK})$ is an immediate consequence of the adversarial bandit lower bound (Theorem 3.5) under the assumption that a constant fraction of the contexts in \mathcal{S} marks at least constant fraction of the n game rounds.

4.2 The Expert Case

We now consider the contextual variant of the basic adversarial bandit model of Section 3. In this variant there is a finite set of N randomized policies. Following the setting of prediction with expert advice, no assumptions are made on the way policies compute their randomized predictions, and the forecaster experiences the contexts only through the advice provided by the policies. For this reason, in what follows we use the word expert to denote a policy. Calling this a model of contextual bandits may sound a little strange, as the structure of contexts does not seem to play a role here. However, we have decided to include this setting in this section because bandit with experts have been used in practical contextual bandit problems — see, e.g., the news recommendation experiment in Ref. [40].

Formally, at each step $t = 1, 2, \dots$ the forecaster obtains the expert advice $(\xi_t^1, \dots, \xi_t^N)$, where each ξ_t^j is a probability distribution over arms representing the randomized play of expert j at time t . If $\ell_t = (\ell_{1,t}, \dots, \ell_{K,t}) \in [0, 1]^K$ is the vector of losses incurred by the K

arms at time t , then $\mathbb{E}_{i \sim \xi_t^j} \ell_{i,t}$ denotes the expected loss of expert j at time t . We allow the expert advice to depend on the realization of the forecaster's past random plays. This fact is explicitly used in the proof of Theorem 4.5.

Similarly to the pseudo-regret (Equation (3.1)) for adversarial bandits, we now introduce the pseudo-regret \bar{R}_n^{ctx} for the adversarial contextual bandit problem,

$$\bar{R}_n^{\text{ctx}} = \max_{i=1,\dots,N} \mathbb{E} \left[\sum_{t=1}^n \ell_{I_t,t} - \sum_{t=1}^n \mathbb{E}_{k \sim \xi_t^i} \ell_{k,t} \right].$$

In order to bound the contextual pseudo-regret \bar{R}_n^{ctx} , one could naively use Exp3 strategy of Section 3 on the set of experts. This would give a bound of order $\sqrt{nN \log N}$. In Figure 4.1 we introduce the contextual forecaster Exp4 for which we show a bound of order $\sqrt{nK \ln N}$. Thus, in this framework we can be competitive even with an exponentially large (with respect to n) number of experts.

Exp4 is a simple adaptation of Exp3 to the contextual setting. Exp4 runs Exp3 over the N experts using estimates of the experts' losses $\mathbb{E}_{i \sim \xi_t^j} \ell_{i,t}$. In order to draw arms, Exp4 mixes the expert advice with the probability distribution over experts maintained by Exp3. The resulting bound on the pseudo-regret is of order $\sqrt{nK \ln N}$, where the term $\sqrt{\ln N}$ comes from running Exp3 over the N experts, while \sqrt{K} is a bound on the second moment of the estimated expert losses under the distribution q_t computed by Exp3. Equation (4.6) shows that $\mathbb{E}_{j \sim q_t} \tilde{y}_{j,t}^2 \leq \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t}^2$. That is, this second moment is at most that of the estimated arm losses under the distribution p_t computed by Exp4, which in turn is bounded by \sqrt{K} using techniques from Section 3.

Theorem 4.2 (Pseudo-regret of Exp4). Exp4 without mixing and with $\eta_t = \eta = \sqrt{\frac{2 \ln N}{nK}}$ satisfies

$$\bar{R}_n^{\text{ctx}} \leq \sqrt{2nN \ln K}. \quad (4.1)$$

On the other hand, with $\eta_t = \sqrt{\frac{\ln N}{tK}}$ it satisfies

$$\bar{R}_n^{\text{ctx}} \leq 2\sqrt{nN \ln K}. \quad (4.2)$$

Exp4 (Exponential weights algorithm for Exploration and Exploitation with Experts) without mixing:

Parameter: a nonincreasing sequence of real numbers $(\eta_t)_{t \in \mathbb{N}}$.

Let q_1 be the uniform distribution over $\{1, \dots, N\}$.

For each round $t = 1, 2, \dots, n$

- (1) Get expert advice ξ_t^1, \dots, ξ_t^N , where each ξ_t^j is a probability distribution over arms.
- (2) Draw an arm I_t from the probability distribution $p_t = (p_{1,t}, \dots, p_{K,t})$, where $p_{i,t} = \mathbb{E}_{j \sim q_t} \xi_{i,t}^j$.
- (3) Compute the estimated loss for each arm

$$\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{p_{i,t}} \mathbb{1}_{I_t=i} \quad i = 1, \dots, K.$$

- (4) Compute the estimated loss for each expert

$$\tilde{y}_{j,t} = \mathbb{E}_{i \sim \xi_t^j} \tilde{\ell}_{i,t} \quad j = 1, \dots, N.$$

- (5) Update the estimated cumulative loss for each expert $\tilde{Y}_{j,t} = \sum_{s=1}^t \tilde{y}_{j,s}$ for $j = 1, \dots, N$.
- (6) Compute the new probability distribution over the experts $q_{t+1} = (q_{1,t+1}, \dots, q_{N,t+1})$, where

$$q_{j,t+1} = \frac{\exp(-\eta_t \tilde{Y}_{j,t})}{\sum_{k=1}^N \exp(-\eta_t \tilde{Y}_{k,t})}.$$

Fig. 4.1 Exp4 forecaster.

Proof. We apply the analysis of Exp3 (Theorem 3.1) to a forecaster using distributions q_t over N experts, whose pseudo-losses are $\tilde{y}_{j,t}$ for $j = 1, \dots, N$. This immediately gives the inequality

$$\sum_{t=1}^n \mathbb{E}_{j \sim q_t} \tilde{y}_{j,t} \leq \tilde{Y}_{k,n} + \frac{\log N}{\eta_n} + \frac{1}{2} \sum_{t=1}^n \eta_t \mathbb{E}_{j \sim q_t} \tilde{y}_{j,t}^2. \quad (4.3)$$

Now, similarly to Equation (3.5) in the proof of Theorem 3.1, we establish the following inequalities

$$\mathbb{E}_{I_t \sim p_t} \tilde{y}_{k,t} = \mathbb{E}_{I_t \sim p_t} \mathbb{E}_{i \sim \xi_t^k} \tilde{\ell}_{i,t} = \mathbb{E}_{i \sim \xi_t^k} \ell_{i,t} = y_{k,t} \quad (4.4)$$

$$\mathbb{E}_{j \sim q_t} \tilde{y}_{j,t} = \mathbb{E}_{j \sim q_t} \mathbb{E}_{i \sim \xi_t^j} \tilde{\ell}_{i,t} = \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t} = \ell_{I_t,t} \quad (4.5)$$

$$\mathbb{E}_{j \sim q_t} \tilde{y}_{j,t}^2 = \mathbb{E}_{j \sim q_t} (\mathbb{E}_{i \sim \xi_t^j} \tilde{\ell}_{i,t})^2 \leq \mathbb{E}_{j \sim q_t} \mathbb{E}_{i \sim \xi_t^j} \tilde{\ell}_{i,t}^2 = \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t}^2 = \frac{\ell_{I_t,t}^2}{p_{I_t,t}} \quad (4.6)$$

where we used Jensen's inequality to prove Equation (4.6). By applying Equations (4.5) and (4.6) to Equation (4.3) we get

$$\sum_{t=1}^n \ell_{I_t,t} = \sum_{t=1}^n \mathbb{E}_{j \sim q_t} \tilde{y}_{j,t} \leq \tilde{Y}_{k,n} + \frac{\log N}{\eta_n} + \frac{1}{2} \sum_{t=1}^n \eta_t \frac{\ell_{I_t,t}^2}{p_{I_t,t}}.$$

Now note that, if we take expectation over the draw of I_1, \dots, I_n , using Equation (4.4) we obtain

$$\mathbb{E} \tilde{Y}_{k,n} = \mathbb{E} \left[\sum_{t=1}^n \mathbb{E}[\tilde{y}_{j,n} | I_1, \dots, I_{t-1}] \right] = \mathbb{E} \left[\sum_{t=1}^n \mathbb{E}_{i \sim \xi_t^k} \ell_{i,t} \right] = \mathbb{E} Y_{k,n}.$$

Hence,

$$\bar{R}_n^{\text{ctx}} = \max_{k=1, \dots, N} \mathbb{E} \left[\sum_{t=1}^n \ell_{I_t,t} - Y_{k,n} \right] \leq \frac{\log N}{\eta_n} + \frac{K}{2} \sum_{t=1}^n \eta_t.$$

Choosing η_t as in the statement of the theorem, and using the inequality $\sum_{t=1}^n t^{-1/2} \leq 2\sqrt{n}$, concludes the proof. \square

Besides pseudo-regret, the contextual regret

$$R_n^{\text{ctx}} = \max_{k=1, \dots, N} \left(\sum_{t=1}^n \ell_{I_t,t} - \sum_{t=1}^n \mathbb{E}_{i \sim \xi_t^k} \ell_{i,t} \right)$$

can be also bounded, at least with high probability. Indeed, similarly to the variant Exp3.P of Exp3 (see Section 3.2), an analogous modification of Exp4, called Exp4.P, satisfies

$$R_n^{\text{ctx}} \leq c\sqrt{nK \ln(N\delta^{-1})}$$

for some constant $c > 0$ and with probability at least $1 - \delta$, where $\delta \in (0, 1)$ is a parameter of the algorithm.

4.2.1 Competing Against the Best Context Set

We revisit the basic contextual scenario introduced in Section 4.1, where the goal is to compete against the best mapping from contexts to arms. Consider now a class $\{\mathcal{S}_\theta : \theta \in \Theta\}$ of context sets. In this new game, each time step $t = 1, 2, \dots$ is marked by the vector $(s_{\theta,t})_{\theta \in \Theta}$ of contexts, one for each set in Θ . Introduce the pseudoregret

$$\bar{R}_n^\Theta = \max_{\theta \in \Theta} \max_{g: \mathcal{S}_\theta \rightarrow \{1, \dots, K\}} \mathbb{E} \left[\sum_{t=1}^n \ell_{I_t, t} - \sum_{t=1}^n \ell_{g(s_{\theta,t}), t} \right].$$

When $|\Theta| = 1$ we recover the contextual pseudoregret $\bar{R}_n^{\mathcal{S}}$. In general, when Θ contains more than one set, the forecaster must learn both the best set \mathcal{S}_θ and the best function $g: \mathcal{S}_\theta \rightarrow \{1, \dots, K\}$ from that set to the set of arms.

We find this variant of contextual bandits interesting because its solution involves a nontrivial combination of two of the main algorithms examined in this section: Exp4 and \mathcal{S} -Exp3. In particular, we consider a scenario in which Exp4 uses instances of \mathcal{S} -Exp3 as experts. The interesting aspect is that these experts are learning themselves, and thus the analysis of the combined algorithm requires taking into account the learning process at both levels.

Note that in order to solve this problem we could simply lump all contexts in a big set and use the proof of Theorem 4.1. However, this would give a regret bound that depends exponentially in $|\Theta|$. On the other hand, by using Exp4 directly on the set of all policies g (which is of cardinality exponential in $|\Theta| \times |S|$), we could improve this to a bound that scales with $\sqrt{|\Theta|}$. The idea we explore here is to use Exp4 over the class Θ of “experts” and combine this with the \mathcal{S} -Exp3 algorithm of Theorem 4.1. This gets us down to a logarithmic dependency on $|\Theta|$, albeit at the price of a worse dependency on n .

Intuitively, Exp4 provides competitiveness against the best context set \mathcal{S}_θ , while the instances of the \mathcal{S} -Exp3 algorithm, acting as experts for Exp4, ensure that we are competitive against the best function $g: \mathcal{S}_\theta \rightarrow \{1, \dots, K\}$ for each $\theta \in \Theta$. However, by doing so we immediately run into a problem: the p_t used by Exp4 is not the same as the p_t 's used by each expert. In order to address this issue, we now show that

the analysis of Exp3 holds even when the sequence of plays I_1, I_2, \dots is drawn from a sequence of distributions q_1, q_2, \dots possibly different from the one chosen by the forecaster. The only requirement we need is that each probability in q_t be bounded away from zero.

Theorem 4.3. Consider a K -armed bandit game in which at each step $t = 1, 2, \dots$ the played arm I_t is drawn from an arbitrary distribution q_t over arms. Each q_t may depend in an arbitrary way on the pairs $(I_1, \ell_{I_1,1}), \dots, (I_{t-1}, \ell_{I_{t-1},t-1})$. Moreover, $q_{t,i} \geq \varepsilon > 0$ for all $i = 1, \dots, K$ and $t \geq 1$.

If Exp3 without mixing is run with $\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{q_{i,t}} \mathbf{1}_{I_t=i}$ and $\eta_t = \eta = \sqrt{\frac{2 \ln K}{nK}}$ then

$$\max_{k=1, \dots, K} \mathbb{E}_{I^n \sim q^n} \left[\sum_{t=1}^n \mathbb{E}_{i \sim p_t} \ell_{i,t} - \sum_{t=1}^n \ell_{k,t} \right] \leq \sqrt{\frac{2n}{\varepsilon} \ln K} \quad (4.7)$$

where $I^n \sim q^n$ means that each I_t is drawn from q_t for $t = 1, \dots, n$, and p_t is the distribution used by Exp3 at time t .

Proof. The proof is an easy adaptation of Exp3 analysis (Theorem 3.1 in Section 3) and we just highlight the differences. The key step is the analysis of the log moment of $\tilde{\ell}_{i,t}$:

$$\begin{aligned} \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t} &= \frac{1}{\eta} \log \mathbb{E}_{i \sim p_t} \exp(-\eta(\tilde{\ell}_{i,t} - \mathbb{E}_{k \sim p_t} \tilde{\ell}_{k,t})) \\ &\quad - \frac{1}{\eta} \log \mathbb{E}_{i \sim p_t} \exp(-\eta \tilde{\ell}_{i,t}). \end{aligned}$$

The first term is bounded in a manner slightly different from the proof of Theorem 3.1,

$$\log \mathbb{E}_{i \sim p_t} \exp(-\eta(\tilde{\ell}_{i,t} - \mathbb{E}_{k \sim p_t} \tilde{\ell}_{k,t})) \leq \frac{\eta^2}{2} \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t}^2 \leq \frac{\eta^2}{2} \frac{p_{I_t,t}}{q_{I_t,t}^2}.$$

The analysis of the second term is unchanged: Let $\tilde{L}_{i,0} = 0$, $\Phi_0(\eta) = 0$ and $\Phi_t(\eta) = \frac{1}{\eta} \log \frac{1}{K} \sum_{i=1}^K \exp(-\eta \tilde{L}_{i,t})$. Then by definition of p_t we have:

$$-\frac{1}{\eta} \log \mathbb{E}_{i \sim p_t} \exp(-\eta \tilde{\ell}_{i,t}) = \Phi_{t-1}(\eta) - \Phi_t(\eta).$$

Proceeding again as in the proof of Theorem 3.1 we obtain

$$\mathbb{E}_{I^n \sim q^n} \left[\sum_{t=1}^n \mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t} \right] \leq \mathbb{E}_{I^n \sim q^n} \left[\sum_{t=1}^n \tilde{\ell}_{k,t} + \frac{\eta}{2} \frac{p_{I_t,t}}{q_{I_t,t}^2} \right] + \frac{\ln K}{\eta}.$$

Now observe that

$$\mathbb{E}_{I_t \sim q_t} \tilde{\ell}_{k,t} = \ell_{k,t} \quad \text{and} \quad \mathbb{E}_{I_t \sim q_t} \frac{p_{I_t,t}}{q_{I_t,t}^2} = \sum_{i=1}^K \frac{p_{i,t}}{q_{i,t}} \leq \frac{1}{\varepsilon}.$$

Therefore,

$$\begin{aligned} \mathbb{E}_{I^n \sim q^n} \left[\sum_{t=1}^n (\mathbb{E}_{i \sim p_t} \ell_{i,t} - \ell_{k,t}) \right] &= \mathbb{E}_{I^n \sim q^n} \left[\sum_{t=1}^n (\mathbb{E}_{i \sim p_t} \tilde{\ell}_{i,t} - \tilde{\ell}_{k,t}) \right] \\ &\leq \frac{\eta n}{2\varepsilon} + \frac{\ln K}{\eta}. \end{aligned}$$

Choosing η as in the statement of the theorem concludes the proof. \square

It is left to the reader to verify that the analysis of \mathcal{S} -Exp in Theorem 4.1 can be combined with the above analysis to give the bound

$$\max_{g: \mathcal{S} \rightarrow \{1, \dots, K\}} \mathbb{E}_{I^n \sim q^n} \left[\sum_{t=1}^n \mathbb{E}_{i \sim p_t} \ell_{i,t} - \sum_{t=1}^n \ell_{g(s_t),t} \right] \leq \sqrt{\frac{2n}{\varepsilon}} |\mathcal{S}| \ln K. \quad (4.8)$$

Next, we state a bound on the contextual pseudoregret of a variant of Exp4 whose probabilities $p_{i,t}$ satisfy the property $p_{i,t} \geq \frac{\gamma}{K}$ for all $i = 1, \dots, K$ and $t \geq 1$, where $\gamma > 0$ is a parameter. This is obtained by replacing in Exp4 the assignment $p_{i,t} = \mathbb{E}_{j \sim q_t} \xi_{i,t}^j$ (line 2 in Figure 4.1) with the assignment

$$p_{i,t} = (1 - \gamma) \mathbb{E}_{j \sim q_t} \xi_{i,t}^j + \frac{\gamma}{K}$$

where $\gamma > 0$ is the mixing coefficient. This mixing clearly achieves the desired property for each $p_{i,t}$.

Theorem 4.4 (Pseudo-regret of Exp4 with mixing). Exp4 with mixing coefficient γ and with $\eta_t = \eta = \gamma/K$ satisfies

$$\overline{R}_n^{\text{ctx}} \leq \frac{\gamma n}{2} + \frac{K \ln N}{\gamma}. \quad (4.9)$$

Proof. The proof goes along the same lines of Exp4 original proof [24, Theorem 7.1] with the following modifications: since the weights are negative exponentials, we can use the bound $\exp(-x) \leq 1 - x + \frac{x^2}{2}$ for all $x \geq 0$ rather than $\exp(x) \leq 1 + x + (e - 2)x^2$ for all $0 \leq x \leq 1$; the term $(1 - \gamma) \sum_t \ell_{k,t}$ is upper bounded directly by $\sum_t \ell_{k,t}$; the term $\frac{\gamma}{K} \sum_t \sum_i \ell_{i,t}$ is upper bounded by γn without requiring the assumption that the expert set contains the “uniform expert”. Finally, the fact that experts’ distributions ξ_t^j depend on the realization of past forecaster’s random arms is dealt with in the same way as in the proof of Theorem 4.2. \square

Theorem 4.5. There exists a randomized forecaster achieving

$$\bar{R}_n^\Theta = \mathcal{O} \left(n^{2/3} \left(\max_{\theta \in \Theta} |\mathcal{S}_\theta| K \ln K \right)^{1/3} \sqrt{\ln |\Theta|} \right)$$

for any class $\{\mathcal{S}_\theta : \theta \in \Theta\}$ of context sets.

Proof. We run the Exp4 forecaster with mixing coefficient γ using instances of the \mathcal{S} -Exp3 algorithm (defined in the proof of Theorem 4.1) as experts. Each \mathcal{S} -Exp3 instance is run on a different context set \mathcal{S}_θ for $\theta \in \Theta$. Let ξ_t^θ be the distribution used at time t by the \mathcal{S} -Exp3 instance running on context set \mathcal{S}_θ and let p^n be the joint distribution of $I^n = (I_1, \dots, I_n)$ used by Exp4. Since $p_{i,t} \geq \frac{\gamma}{K}$ for all $i = 1, \dots, K$ and $t \geq 1$, we can use Equation (4.8) with $\varepsilon = \gamma/K$. Thus, Theorem 4.4 implies

$$\begin{aligned} \mathbb{E}_{I^n \sim p^n} \left[\sum_{t=1}^n \ell_{I_t,t} \right] &\leq \min_{\theta \in \Theta} \mathbb{E}_{I^n \sim p^n} \left[\sum_{t=1}^n \mathbb{E}_{k \sim \xi_t^\theta} \ell_{k,t} \right] + \frac{\gamma n}{2} + \frac{K \ln |\Theta|}{\gamma} \\ &\leq \min_{\theta \in \Theta} \min_{g: \mathcal{S}_\theta \rightarrow \{1, \dots, K\}} \mathbb{E} \left[\sum_{t=1}^n \ell_{g(s_t),t} \right] \\ &\quad + \sqrt{\frac{2n}{\varepsilon} \max_{\theta \in \Theta} |\mathcal{S}_\theta| \ln K} + \frac{\gamma n}{2} + \frac{K \ln |\Theta|}{\gamma}. \end{aligned}$$

Substituting $\varepsilon = \gamma/K$ in the above expression and choosing γ of the order of $n^{-1/3} (\max_{\theta \in \Theta} |\mathcal{S}_\theta| K \ln K)^{1/3} \sqrt{\ln |\Theta|}$ gives the desired result. \square

Note that in Theorem 4.5 the rate is $n^{2/3}$, in contrast to the more usual $n^{1/2}$ bandit rate. This worsening is inherent in the Exp4-over-Exp3 construction. It is not known whether the rate could be improved while keeping the same logarithmic dependence on $|\Theta|$ guaranteed by this construction.

4.3 Stochastic Contextual Bandits

We now move on to consider the case in which policies have a known structure. More specifically, each policy is a function f mapping the context space to the arm space $\{1, \dots, K\}$ and the set \mathcal{F} of policies is given as an input parameter to the forecaster.

Under this assumption on the policies, the problem can be viewed as a bandit variant of supervised learning. For this reason, here and in the next section we follow the standard notation of supervised learning and use x rather than s to denote contexts.

In supervised learning, we observe data of the form (x_t, ℓ_t) . In the contextual bandit setting, the observed data are $(x_t, \ell_{I_t, t})$ where I_t is the arm chosen by the forecaster at time t given context $x_t \in \mathcal{X}$. This connection to supervised learning has steered the focus of research toward stochastic data generation models, which are widespread in the analysis of supervised learning. In the stochastic variant of contextual bandits, contexts x_t and arm losses $\ell_t = (\ell_{1,t}, \dots, \ell_{K,t})$ are realizations of i.i.d. draws from a fixed and unknown distribution D over $\mathcal{X} \times [0, 1]^K$. In tight analogy with statistical learning theory, a policy f is evaluated in terms of its statistical risk $\ell_D(f) = \mathbb{E}_{(x, \ell) \sim D} \ell_{f(x)}$. Let

$$f^* = \operatorname{arginf}_{f \in \mathcal{F}} \ell_D(f)$$

the risk-minimizing policy in the class. The regret with respect to the class \mathcal{F} of a forecaster choosing arms I_1, I_2, \dots is then defined by

$$\sum_{t=1}^n \ell_{I_t, t} - n \ell_D(f^*).$$

This can be viewed as the stochastic counterpart of the adversarial contextual regret \bar{R}_n^{ctx} introduced in Section 4.2. The main question is

now to characterize the “price of bandit information” using the sample complexity of supervised learning as yardstick.

In the rest of this section we focus on the case of $K = 2$ arms and parameterize classes \mathcal{F} of policies $f : \mathcal{X} \rightarrow \{1, 2\}$ by their VC-dimension d — see [41] for a modern introduction to VC theory. For this setting, we consider the following forecaster.

VE (VC dimension by Exponentiation):

Parameters: number n of rounds, n' satisfying $1 \leq n' \leq n$.

- (1) For the first n' rounds, choose arms uniformly at random.
- (2) Build $\mathcal{F}' \subseteq \mathcal{F}$ such that for any $f \in \mathcal{F}$ there is exactly one $f' \in \mathcal{F}'$ satisfying $f(x_t) = f'(x_t)$ for all $t = 1, \dots, n'$.
- (3) For $t = n' + 1, \dots, n$ play by simulating Exp4.P using the policies of \mathcal{F}' as experts.

We now show that the per round regret of VE is of order $\sqrt{d/n}$, excluding logarithmic factors. This rate is equal to the optimal rate for supervised learning of VC-classes, showing that — in this case — the price of bandit information is essentially zero.

Theorem 4.6. For any class \mathcal{F} of binary policies $f : \mathcal{X} \rightarrow \{0, 1\}$ of VC-dimension d and for all $n > d$, the forecaster VE run with $n' = \sqrt{n(2d \ln \frac{en}{d} + \ln \frac{3}{\delta})}$ satisfies

$$\sum_{t=1}^n \ell_{I_t, t} - n \inf_{f \in \mathcal{F}} \ell_D(f) \leq c \sqrt{n \left(d \ln \frac{en}{d} + \ln \frac{3}{\delta} \right)} \quad (4.10)$$

for some constant $c > 0$ and with probability at least $1 - \delta$ with respect to both the random data generation and VE’s internal randomization.

Proof. Given a sample realization $(x_1, \ell_1), \dots, (x_n, \ell_n)$, let f' be the unique element of \mathcal{F}' such that $f'(x_t) = f^*(x_t)$ for all $t = 1, \dots, n'$, where f^* is the risk-minimizing function in \mathcal{F} . Given a sample, without

loss of generality we may assume that \mathcal{F} contains functions restricted on the finite domain $\{x_1, \dots, x_n\}$. Recall Sauer–Shelah lemma — see, e.g., [41] — stating that any class \mathcal{F} of binary functions defined on a finite domain of size n satisfies $|\mathcal{F}| \leq (\frac{en}{d})^d$, where d is the VC-dimension of \mathcal{F} . Then, with probability at least $1 - \frac{\delta}{3}$ with respect to VE’s internal randomization,

$$\begin{aligned}
\sum_{t=1}^n \ell_{I_t, t} &\leq n' + \sum_{t=n'+1}^n \ell_{f'(x_t), t} + c\sqrt{2(n-n') \ln \frac{3|\mathcal{F}'|}{\delta}} \\
&\leq n' + \sum_{t=n'+1}^n (\ell_{f^*(x_t), t} + \ell_{f'(x_t), t} - \ell_{f^*(x_t), t}) \\
&\quad + c\sqrt{2(n-n') \ln \frac{3|\mathcal{F}'|}{\delta}} \\
&\leq n' + \sum_{t=n'+1}^n (\ell_{f^*(x_t), t} + \mathbb{1}_{f'(x_t) \neq f^*(x_t)}) + c\sqrt{2(n-n') \ln \frac{3|\mathcal{F}'|}{\delta}} \\
&\leq n' + \sum_{t=n'+1}^n (\ell_{f^*(x_t), t} + \mathbb{1}_{f'(x_t) \neq f^*(x_t)}) \\
&\quad + c\sqrt{2n \left(d \ln \frac{en}{d} + \ln \frac{3}{\delta} \right)}
\end{aligned}$$

where we used $\ell_{i,t} \in [0, 1]$ in the penultimate step and the Sauer–Shelah lemma in the last step. Now, the term $\sum_t \ell_{f^*(x_t), t}$ is controlled in probability w.r.t. the random draw of the sample via Chernoff bounds,

$$\mathbb{P} \left(\sum_{t=n'+1}^n \ell_{f^*(x_t), t} > (n-n') \ell_D(f^*) + \sqrt{\frac{n-n'}{2} \ln \frac{3}{\delta}} \right) \leq \delta.$$

Hence,

$$\begin{aligned}
\sum_{t=1}^n \ell_{I_t, t} &\leq n' + n \ell_D(f^*) \\
&\quad + \sum_{t=n'+1}^n \mathbb{1}_{f'(x_t) \neq f^*(x_t)} + c\sqrt{2n \left(d \ln \frac{en}{d} + \ln \frac{3}{\delta} \right)}
\end{aligned}$$

with probability at least $\frac{2\delta}{3}$ with respect to both the random sample draw and VE's internal randomization.

The term $\sum_t \mathbb{1}_{f'(x_t) \neq f^*(x_t)}$ quantifies the fact that the unique function $f' \in \mathcal{F}'$ that agrees with f^* on the first n' data points is generally different from f^* on the remaining $n - n'$ points. Since each data point (x_t, ℓ_t) is drawn i.i.d., the distribution of a sequence of n pairs remains the same if we randomly permute their positions after drawing them. Hence we can bound $\sum_t \mathbb{1}_{f'(x_t) \neq f^*(x_t)}$ in probability w.r.t. a random permutation σ of $\{1, \dots, n\}$. Let $\|f - g\| = \sum_{t=1}^n \mathbb{1}_{f'(x_t) \neq f^*(x_t)}$. Then

$$\begin{aligned} & \mathbb{P}_\sigma \left(\sum_{t=n'+1}^n \mathbb{1}_{f'(x_{\sigma(t)}) \neq f^*(x_{\sigma(t)})} > k \right) \\ & \leq \mathbb{P}_\sigma \left(\exists f, g \in \mathcal{F}, \|f - g\| > k : f(x_{\sigma(t)}) = g(x_{\sigma(t)}), t = 1, \dots, n' \right) \\ & \leq |\mathcal{F}|^2 \left(1 - \frac{k}{n} \right)^{n'} \\ & \leq \left(\frac{en}{d} \right)^{2d} \exp \left(-\frac{kn'}{n} \right) \\ & \leq \frac{\delta}{3} \end{aligned}$$

for

$$k \geq \frac{n}{n'} \left(2d \ln \frac{en}{d} + \ln \frac{3}{\delta} \right).$$

Now, since we just proved that

$$\sum_{t=n'+1}^n \mathbb{1}_{f'(x_{\sigma(t)}) \neq f^*(x_{\sigma(t)})} \leq \frac{n}{n'} \left(2d \ln \frac{en}{d} + \ln \frac{3}{\delta} \right)$$

holds with probability at least $\frac{\delta}{3}$ for any sample realization, it holds with the same probability for a random sample. Hence, by choosing n' as in the statement of the theorem and overapproximating, we get the desired result. \square

4.4 The Multiclass Case

A different viewpoint on contextual bandits is provided by the so-called bandit multiclass problem. This is a bandit variant of the online

protocol for multiclass classification, where the goal is to sequentially learn a mapping from the context space \mathbb{R}^d to the label space $\{1, \dots, K\}$, with $K \geq 2$. In this protocol the learner keeps a classifier parameterized by a $K \times d$ matrix W . At each time step $t = 1, 2, \dots$ the side information $x_t \in \mathbb{R}^d$ is observed (following standard notations in online classification, here we use x instead of s to denote contexts), and the learner predicts the label \hat{y}_t maximizing $(Wx_t)_i$ over all labels $i = 1, \dots, K$. In the standard online protocol, the learner observes the true label y_t associated with x_t after each prediction and uses this information to adjust W . In the bandit version, the learner only observes $\mathbb{1}_{\hat{y}_t \neq y_t}$, that is, whether the prediction at time t was correct or not.

A simple but effective learning strategy for (non-bandit) online classification is the multiclass Perceptron algorithm. This algorithm updates W at time t using the rule $W \leftarrow W + X_t$, where X_t is a $K \times d$ matrix with components $(X_t)_{i,j} = x_{t,j}(\mathbb{1}_{y_t=i} - \mathbb{1}_{\hat{y}_t=i})$. Therefore, the update rule can be rewritten as

$$\begin{aligned} w_{y_t} &\leftarrow w_{y_t} + x_t \\ w_{\hat{y}_t} &\leftarrow w_{\hat{y}_t} - x_t \\ w_i &\leftarrow w_i \quad \text{for all } i \neq y_t \text{ and } i \neq \hat{y}_t \end{aligned}$$

where w_i denotes the i -th row of matrix W . Note, in particular, that no update takes place (i.e., X_t is the all zero matrix) when $\hat{y}_t = y_t$, which means that y_t is predicted correctly.

A straightforward generalization of the Perceptron analysis gives that, on any sequence of $(x_1, y_1), (x_2, y_2), \dots \in \mathbb{R}^d \times \{1, \dots, K\}$ such that $\|x_t\| = 1$, the number of classification mistakes satisfies the following notion of regret,

$$\sum_{t=1}^n \mathbb{1}_{\hat{y}_t \neq y_t} \leq \inf_U (L_n(U) + 2\|U\|^2 + \|U\| \sqrt{2n\bar{L}_n(U)})$$

uniformly over $n \geq 1$, where the infimum is over all $K \times d$ matrices U and $\|\cdot\|$ denotes the Frobenius norm. Here $L_n(U)$ denotes the cumulative hinge loss of policy U ,

$$L_n(U) = \sum_{t=1}^n \ell_t(U) = \sum_{t=1}^n \left[1 - (Ux_t)_{y_t} + \max_{i \neq y_t} (Ux_t)_i \right]_+$$

where $[\cdot]_+ = \max\{0, \cdot\}$ is the *hinge* function. Finally, $\bar{L}_n(U) = \frac{1}{n}L_n(U)$ is the average hinge loss of U .

Note that $\ell_t(U) = 0$ if and only if $(Ux_t)_{y_t} \geq 1 + \max_{i \neq y_t} (Ux_t)_i$, which can only happen when $y_t = \hat{y}_t = \operatorname{argmax}_i (Ux_t)_i$. Moreover, $\ell_t(U) \geq 1$ if and only if $\hat{y}_t \neq y_t$. This means that ℓ_t is a convex upper bound on the mistake indicator function $\mathbb{1}_{\hat{y}_t \neq y_t}$ for the multiclass classifier represented by U .

We now introduce a bandit variant of the multiclass Perceptron called Banditron.

Banditron

Parameter: number $\gamma \in (0, \frac{1}{2})$.

Initialize: Set W_1 to the zero $K \times d$ matrix.

For each round $t = 1, 2, \dots, n$

- (1) Observe $x_t \in \mathbb{R}^d$.
- (2) Set $\hat{y}_t = \operatorname{argmax}_{i=1, \dots, K} (W_t x_t)_i$.
- (3) Predict $Y_t \in \{1, \dots, K\}$ drawn from distribution $p_t = (p_{1,t}, \dots, p_{K,t})$ such that $p_{i,t} = (1 - \gamma)\mathbb{1}_{\hat{y}_t=i} + \frac{\gamma}{K}$.
- (4) Observe $\mathbb{1}_{Y_t=y_t}$.
- (5) Update $W_{t+1} = W_t + \tilde{X}_t$ where

$$(\tilde{X}_t)_{i,j} = x_{t,j} \left(\frac{\mathbb{1}_{Y_t=y_t} \mathbb{1}_{Y_t=i}}{p_{i,t}} - \mathbb{1}_{\hat{y}_t=i} \right).$$

The Banditron operates in the bandit variant of the online protocol for multiclass classification. As X_t depends on the true label y_t , which is only available when the classification is correct, the Banditron computes an estimate \tilde{X}_t of X_t via a randomized technique inspired by the nonstochastic multi-armed bandit problem. The label $\hat{y}_t = \operatorname{argmax}_i (W x_t)_i$ is used to make the prediction at time t only with probability $1 - \gamma$, whereas with probability γ a random label is predicted at each time t .

We now analyze the expected number of prediction mistakes made by the Banditron algorithm on any sequence of examples (x_t, y_t) . Unlike the non-bandit case, where the number of mistakes M_n after n steps of

the multiclass Perceptron provides a “multiclass regret” bound $M_n - L_n(U) = \mathcal{O}(\sqrt{n})$, in the bandit case the regret achieved by the variant of the Perceptron is only bounded by $\mathcal{O}(n^{2/3})$.

Theorem 4.7. If the Banditron algorithm is run with parameter $\gamma = (K/n)^{1/3}$ on any sequence $(x_1, y_1), \dots, (x_n, y_n) \in \mathbb{R}^d \times \{-1, +1\}$ of examples such that $n \geq 8K$ and $\|x_t\| = 1$, then the number M_n of prediction mistakes satisfies

$$\begin{aligned} \mathbb{E} M_n &\leq \inf_U (L_n(U) + (1 + \|U\| \sqrt{2\bar{L}_n(U)}) K^{1/3} n^{2/3} \\ &\quad + 2\|U\|^2 K^{2/3} n^{1/3} + \sqrt{2}\|U\| K^{1/6} n^{1/3}) \end{aligned}$$

where the infimum is over all $K \times d$ matrices U and $\bar{L}_n(U) = \frac{1}{n} L_n(U)$ is the average hinge loss of U .

Proof. We need to bound $M = \sum_t \mathbb{1}_{Y_t \neq y_t}$. Let \mathbb{E}_t be the expectation conditioned on the first $t - 1$ predictions. We start by bounding the first and second moments of \tilde{X}_t ,

$$\begin{aligned} \mathbb{E}_t[(\tilde{X}_t)_{i,j}] &= x_{t,j} \sum_{k=1}^K p_{k,t} \left(\frac{\mathbb{1}_{k=y_t} \mathbb{1}_{k=i}}{p_{k,t}} - \mathbb{1}_{\hat{y}_t=i} \right) \\ &= x_{t,j} (\mathbb{1}_{y_t=i} - \mathbb{1}_{\hat{y}_t=i}) = (X_t)_{i,j}. \end{aligned}$$

For the second moment, note that

$$\begin{aligned} \|\tilde{X}_t\|^2 &= \sum_{i=1}^K \sum_{j=1}^d x_{t,j}^2 \left(\frac{\mathbb{1}_{Y_t=y_t} \mathbb{1}_{Y_t=i}}{p_{i,t}} - \mathbb{1}_{\hat{y}_t=i} \right)^2 \\ &= \sum_{i=1}^K \left(\frac{\mathbb{1}_{Y_t=y_t} \mathbb{1}_{Y_t=i}}{p_{i,t}} - \mathbb{1}_{\hat{y}_t=i} \right)^2 \end{aligned}$$

where

$$\sum_{i=1}^K \left(\frac{\mathbb{1}_{Y_t=y_t} \mathbb{1}_{Y_t=i}}{p_{i,t}} - \mathbb{1}_{\hat{y}_t=i} \right)^2 = \begin{cases} \left(\frac{1}{p_{y_t,t}^2} + 1 \right) & \text{if } Y_t = y_t \neq \hat{y}_t \\ \left(\frac{1}{p_{y_t,t}} - 1 \right)^2 & \text{if } Y_t = y_t = \hat{y}_t \\ 1 & \text{otherwise.} \end{cases}$$

Therefore, if $y_t \neq \hat{y}_t$,

$$\begin{aligned}\mathbb{E}_t \|\tilde{X}_t\|^2 &= p_{y_t,t} \left(\frac{1}{p_{y_t,t}^2} + 1 \right) + (1 - p_{y_t,t}) \\ &= 1 + \frac{1}{p_{y_t,t}} = 1 + \frac{K}{\gamma} \leq \frac{2K}{\gamma}\end{aligned}$$

because $p_{i,t} = \gamma$ when $y_t \neq \hat{y}_t$. Otherwise, if $y_t = \hat{y}_t$

$$\begin{aligned}\mathbb{E}_t \|\tilde{X}_t\|^2 &= p_{y_t,t} \left(\frac{1}{p_{y_t,t}} - 1 \right)^2 + (1 - p_{y_t,t}) \\ &= \frac{1}{p_{y_t,t}} - 1 = \frac{1}{1 - \gamma} - 1 \leq 2\gamma\end{aligned}$$

because $p_{i,t} = 1 - \gamma$ when $y_t = \hat{y}_t$ and $\gamma \leq \frac{1}{2}$. Hence,

$$\mathbb{E}_t \|\tilde{X}_t\|^2 \leq 2 \left(\frac{K}{\gamma} \mathbb{1}_{y_t \neq \hat{y}_t} + \gamma \mathbb{1}_{y_t = \hat{y}_t} \right).$$

We are now ready to prove a bound on the expected number of mistakes. Following the standard analysis for the Perceptron algorithm, we derive upper and lower bounds on the expectation of the quantity $\langle U, W_{n+1} \rangle = \text{tr}(U W_{n+1}^\top)$, for an arbitrary $K \times d$ matrix U . First, using Cauchy–Schwartz and Jensen inequalities we obtain

$$\mathbb{E} \langle U, W_{n+1} \rangle \leq \sqrt{\|U\|^2 \mathbb{E} \|W_{n+1}\|^2}.$$

Now

$$\begin{aligned}\mathbb{E}_n [\|W_{n+1}\|^2] &= \mathbb{E}_n [\|W_n\|^2 + 2\langle W_n, \tilde{X}_n \rangle + \|\tilde{X}_n\|^2] \\ &\leq \|W_n\|^2 + \mathbb{E}_n \|\tilde{X}_n\|^2.\end{aligned}$$

In order to see why the inequality holds, note that

$$\begin{aligned}\mathbb{E}_n \langle W_n, \tilde{X}_n \rangle &= \langle W_n, X_n \rangle = \sum_{i=1}^K (W_n x_t)_i (\mathbb{1}_{y_n=i} - \mathbb{1}_{\hat{y}_n=i}) \\ &= (W_n x_n)_{y_n} - (W_n x_n)_{\hat{y}_n} \leq 0\end{aligned}$$

because $\hat{y}_n = \operatorname{argmax}_{i=1,\dots,K} (W_n x_n)_i$ by definition. Therefore, since W_1 is the zero matrix,

$$\begin{aligned} \mathbb{E} \|W_{n+1}\|^2 &\leq \sum_{t=1}^n \mathbb{E} \|\tilde{X}_n\|^2 \\ &\leq 2 \sum_{t=1}^n \left(\frac{K}{\gamma} \mathbb{P}(y_t \neq \hat{y}_t) + \gamma \mathbb{P}(y_t = \hat{y}_t) \right) \\ &\leq \frac{2K}{\gamma} \sum_{t=1}^n \mathbb{P}(y_t \neq \hat{y}_t) + 2\gamma n. \end{aligned}$$

Thus we have

$$\mathbb{E} \langle U, W_{n+1} \rangle \leq \|U\| \sqrt{\frac{2K}{\gamma} \sum_{t=1}^n \mathbb{P}(y_t \neq \hat{y}_t) + 2\gamma n}.$$

Now we lower bound $\langle U, W_{n+1} \rangle$ as follows,

$$\begin{aligned} \mathbb{E}_n \langle U, W_{n+1} \rangle &= \mathbb{E}_n \langle U, W_n + \tilde{X}_n \rangle \\ &= \langle U, W_n \rangle + \langle U, X_n \rangle \\ &\geq \langle U, W_n \rangle + \mathbf{1}_{y_t \neq \hat{y}_t} - \ell_t(U) \end{aligned}$$

because, by definition of ℓ_t ,

$$\begin{aligned} \ell_t(U) &= \left[1 - (Ux_t)_{y_t} + \max_{i \neq y_t} (Ux_t)_i \right]_+ \\ &\geq 1 - (Ux_t)_{y_t} + (Ux_t)_{\hat{y}_t} \\ &\geq \mathbf{1}_{y_t \neq \hat{y}_t} - (Ux_t)_{y_t} + (Ux_t)_{\hat{y}_t} \\ &= \mathbf{1}_{y_t \neq \hat{y}_t} - \langle U, X_t \rangle. \end{aligned}$$

Therefore, using again the fact that W_1 is the zero matrix,

$$\mathbb{E} \langle U, W_{n+1} \rangle \geq \sum_{t=1}^n \mathbb{P}(y_t \neq \hat{y}_t) - \sum_{t=1}^n \ell_t(U).$$

Combining the upper and lower bounds on $\langle U, W_{n+1} \rangle$ we get

$$\sum_{t=1}^n \mathbb{P}(y_t \neq \hat{y}_t) - L_n(U) \leq \|U\| \sqrt{\frac{2K}{\gamma} \sum_{t=1}^n \mathbb{P}(y_t \neq \hat{y}_t) + 2\gamma n}.$$

Solving for $\sum_t \mathbb{P}(y_t \neq \hat{y}_t)$ and overapproximating yields

$$\begin{aligned} \sum_{t=1}^n \mathbb{P}(y_t \neq \hat{y}_t) &\leq L_n(U) + \frac{2K}{\gamma} \|U\|^2 + \|U\| \sqrt{\frac{2K}{\gamma} L_n(U) + 2\gamma n} \\ &= L_n(U) + \frac{2K}{\gamma} \|U\|^2 + \|U\| \sqrt{\left(\frac{2K}{\gamma} \bar{L}_n(U) + 2\gamma\right) n}. \end{aligned}$$

Now, since $\mathbb{P}(y_t \neq Y_t) = (1 - \gamma)\mathbb{P}(y_t \neq \hat{y}_t) + \gamma$,

$$\begin{aligned} \sum_{t=1}^n \mathbb{P}(y_t \neq Y_t) &\leq L_n(U) + \gamma n + \frac{2K}{\gamma} \|U\|^2 \\ &\quad + \|U\| \sqrt{\left(\frac{2K}{\gamma} \bar{L}_n(U) + 2\gamma\right) n}. \end{aligned}$$

Choosing γ as in the statement of the theorem yields the desired result. Note that $\gamma \leq \frac{1}{2}$ because we assume $n \geq 8K$. \square

4.5 Bibliographic Remarks

A model of contextual stochastic bandits close to those studied here is introduced by Wang et al. [166]. The context is provided by an i.i.d. sequence of random variables and the rewards for each arm depend on the context through an unknown parametric model belonging to a known class. This result has been generalized to the non-i.i.d. case by Wang et al. [165], to the multivariate linear case by Rusmevichientong and Tsitsiklis [148], and to the multivariate and nonparametric case by Perchet and Rigollet [143]. In Ref. [152], contexts and rewards belong to arbitrary metric spaces, and the unknown function mapping contexts to rewards satisfies a Lipschitz assumption (remarkably, the same algorithm also applies to slowly changing expected rewards and sleeping bandit settings). The case of deterministic covariates (fixed design), finitely many arms, and a linear stochastic dependence between covariates and rewards has been studied in [22, 66] — see also [2]. The work of Filippi et al. [78] extends the analysis of fixed design by assuming a generalized linear model to capture the dependence of rewards on covariates.

The simple contextual model analyzed in Section 4.1, as well as its extension described in Section 4.2.1, are due to Maillard and Munos [131]. The Exp4 algorithm for the adversarial case was introduced in Auer et al. [24]. Subsequent improvements were proposed in the two papers, i.e., Beygelzimer et al. [39] (Exp4.P with high-probability bounds) and McMahan and Streeter [136] (exploitation of correlations in expert advice). The VE algorithm and its analysis in Section 4.3 are also taken from Beygelzimer et al. [39].

A stochastic model for contextual bandits with finitely many arms and finitely many states has been investigated by Seldin et al. [150] using new sophisticated tools of PAC-Bayesian analysis.

The general stochastic model of Section 4.3 for contextual bandits with finitely many arms is due to Langford and Zhang [126]. An efficient algorithm for this model has been recently proposed in [74].

The bandit multiclass model of Section 4.4 is due to Langford and Zhang [126]. The Banditron algorithm and its analysis are from Kakade et al. [109]. See also [70, 97] for recent variants and improvements.

5

Linear Bandits

We now consider the important generalization of adversarial bandits where the set of arms $\{1, \dots, K\}$ is replaced by a compact set $\mathcal{K} \subset \mathbb{R}^d$. In this case, the loss at each step is some function defined on \mathcal{K} , and the task is to pick an arm as close as possible to the minimum of the loss function at hand. In order to allow sublinear regret bounds, even in the presence of infinitely many arms, we must assume some structure for the loss function. In particular, in this section we assume that the loss at each time step is a linear function of arms. Linearity is a standard assumption (think, for instance, of linear regression) and naturally occurs in many bandit applications. The source routing problem mentioned in the Introduction is a good example, since the cost of choosing a routing path is linear in the cost of the edges that make up the path. This defines the so-called online linear optimization setting: at each time step $t = 1, 2, \dots$ the forecaster chooses $x_t \in \mathcal{K}$ while, simultaneously, the adversary chooses ℓ_t from some fixed and known subset \mathcal{L} of \mathbb{R}^d . The loss incurred by the forecaster is the inner product $x_t^\top \ell_t$.

In this section we focus on the analysis of the pseudo-regret

$$\bar{R}_n = \mathbb{E} \sum_{t=1}^n x_t^\top \ell_t - \min_{x \in \mathcal{K}} \mathbb{E} \sum_{t=1}^n x^\top \ell_t$$

where the expectation is with respect to the forecaster's internal randomization. The adversarial bandit setting of Section 3 is obtained by choosing $\mathcal{K} = \{e_1, \dots, e_d\}$, where e_1, \dots, e_d is the canonical basis of \mathbb{R}^d , and $\mathcal{L} = [0, 1]^d$. Similarly to Section 3, we focus on the bandit feedback where the forecaster only observes the incurred loss $x_t^\top \ell_t$ at the end round t . However, we also discuss the full information setting, where the complete loss vector ℓ_t is revealed at the end of each round t , as well as other feedback models.

It is important to note that, without any loss of generality (as far as regret bounds are concerned), one can always assume that \mathcal{K} has size $\mathcal{O}(n^d)$. Indeed, since \mathcal{K} is a compact set and the loss is linear (and therefore Lipschitz), one can cover \mathcal{K} with $\mathcal{O}(n^d)$ points such that one incurs a vanishing extra cumulative regret by playing on the discretization rather than on the original set. Of course, the downside of this approach is that a strategy resulting from such a cover is often not computationally efficient. On the other hand, this assumption allows us to apply ideas from Section 3 to this more general setting. We analyze the pseudo-regret for finite classes in Section 5.1. Without loss of generality, it is also possible to assume that \mathcal{K} is a convex set. Indeed, the pseudo-regret is the same if one plays x_t , or if one plays a point at random in \mathcal{K} such that the expectation of the played point is x_t . This remark is critical and allows us to develop a powerful technology known as the Mirror Descent algorithm. We describe this approach in Section 5.2 and explore it further in subsequent sections.

5.1 Exp2 (Expanded Exp) with John's Exploration

In this section we work under the *bounded scalar loss* assumption. That is, we assume that \mathcal{K} and \mathcal{L} are such that $|x^\top \ell| \leq 1$ for any $(x, \ell) \in \mathcal{K} \times \mathcal{L}$. Moreover, we restrict our attention to finite sets \mathcal{K} , with $|\mathcal{K}| = N$. Without loss of generality we assume that \mathcal{K} spans \mathbb{R}^d . If it is not the case, then one can rewrite the elements of \mathcal{K} in some lower dimensional

vector space, and work there. Note that a trivial application of Exp3 to the set \mathcal{K} of arms gives a bound that scales as $\sqrt{nN \ln N}$. If \mathcal{K} is a discretized set (in the sense described earlier), then N is exponential in d . We show here that, by appropriately modifying Exp3, one can obtain a polynomial regret of order $\sqrt{nd \ln N}$.

To describe the strategy, we first need a useful result from convex geometry: John’s theorem. This result concerns the ellipsoid \mathcal{E} of minimal volume enclosing a given convex set \mathcal{K} (which we call John’s ellipsoid of \mathcal{K}). Basically, the theorem states that \mathcal{E} has many contact points with \mathcal{K} , and these contact points are “nicely” distributed, in the sense that they almost form an orthonormal basis — see [32] for a proof.

Theorem 5.1 (John’s theorem). Let $\mathcal{K} \subset \mathbb{R}^d$ be a convex set. If the ellipsoid \mathcal{E} of minimal volume enclosing \mathcal{K} is the unit ball in some norm derived from a scalar product $\langle \cdot, \cdot \rangle$, then there exist $M \leq \frac{1}{2}d(d+1) + 1$ contact points u_1, \dots, u_M between \mathcal{E} and \mathcal{K} , and a probability distribution (μ_1, \dots, μ_M) over these contact points, such that

$$x = d \sum_{i=1}^M \mu_i \langle x, u_i \rangle u_i \quad \forall x \in \mathbb{R}^d.$$

In fact John’s theorem is an *if and only if*, but here we only need the direction stated in the theorem. We are now in position to describe the strategy. Let $\text{Conv}(S)$ be the convex hull of a set $S \in \mathbb{R}^d$. First, we perform a preprocessing step in which the set \mathcal{K} is rewritten so that John’s ellipsoid of $\text{Conv}(\mathcal{K})$ is the unit ball for some inner product $\langle \cdot, \cdot \rangle$. The loss of playing $x \in \mathcal{K}$ against $\ell \in \mathcal{L}$ is then given by $\langle x, \ell \rangle$. See [45] for the details of this transformation. Let $u_1, \dots, u_M \in \mathcal{K}$ and (μ_1, \dots, μ_M) satisfy Theorem 5.1 for the convex set $\text{Conv}(\mathcal{K})$.

Recall from Section 3 that the key idea to play against an adversary is to select x_t at random from some probability distribution p_t over \mathcal{K} . We first describe how to build an unbiased estimate of ℓ_t , given such a point x_t played at random from p_t (such that $p_t(x) > 0$ for any $x \in \mathcal{K}$). Recall that the outer product $x \otimes x$ is defined as the linear mapping from \mathbb{R}^d to \mathbb{R}^d such that $x \otimes x(y) = \langle x, y \rangle x$. Note that one can also view

$x \otimes x$ as a $d \times d$ matrix, so that the evaluation of $x \otimes x$ is equivalent to a multiplication by the corresponding matrix. Now let

$$P_t = \sum_{x \in \mathcal{K}} p_t(x) x \otimes x.$$

Note that this matrix is invertible, since \mathcal{K} is full rank and $p_t(x) > 0$ for all $x \in \mathcal{K}$. The estimate for ℓ_t is given by $\tilde{\ell}_t = P_t^{-1}(x_t \otimes x_t)\ell_t$. Note that this is a valid estimate since $(x_t \otimes x_t)\ell_t = \langle x_t, \ell_t \rangle x_t$ and P_t^{-1} are observed quantities. Also, it is clearly an unbiased estimate.

Now the Exp2 algorithm with John's exploration corresponds to playing according to the following probability distribution

$$p_t(x) = (1 - \gamma) \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \langle x, \tilde{\ell}_s \rangle\right)}{\sum_{y \in \mathcal{K}} \exp\left(-\eta \sum_{s=1}^{t-1} \langle y, \tilde{\ell}_s \rangle\right)} + \gamma \sum_{i=1}^M \mu_i \mathbb{1}_{x=u_i} \quad (5.1)$$

where $\eta, \gamma > 0$ are input parameters. Note that this corresponds to a variant of Exp3 using $\langle x, \tilde{\ell}_t \rangle$ as loss estimate for $x \in \mathcal{K}$, and an exploration distribution supported by the contact points.

Theorem 5.2 (Pseudo-regret of Exp2 with John's exploration).

For any $\eta, \gamma > 0$ such that $\eta d \leq \gamma$, Exp2 with John's exploration satisfies

$$\bar{R}_n \leq 2\gamma n + \frac{\ln N}{\eta} + \eta n d.$$

In particular, with $\eta = \sqrt{\frac{\ln N}{3nd}}$ and $\gamma = \eta d$,

$$\bar{R}_n \leq 2\sqrt{3nd \ln N}.$$

Proof. Since \mathcal{K} is finite, we can easily adapt the analysis of Exp3 in Theorem 3.1 to take into account the exploration term. This gives

$$\bar{R}_n \leq 2\gamma n + \frac{\ln N}{\eta} + \eta \mathbb{E} \sum_{t=1}^n \sum_{x \in \mathcal{K}} p_t(x) \langle x, \tilde{\ell}_t \rangle^2$$

whenever $\eta \langle x, \tilde{\ell}_t \rangle \leq 1$ for all $x \in \mathcal{K}$. We now bound the last term in the right-hand side of the above inequality. Using the definition of the estimated loss $\tilde{\ell}_t = P_t^{-1}(x_t \otimes x_t)\ell_t$, we can write

$$\begin{aligned} \sum_{x \in \mathcal{K}} p_t(x) \langle x, \tilde{\ell}_t \rangle^2 &= \sum_{x \in \mathcal{K}} p_t(x) \langle \tilde{\ell}_t, (x \otimes x) \tilde{\ell}_t \rangle \\ &= \langle \tilde{\ell}_t, P_t \tilde{\ell}_t \rangle \\ &= \langle x_t, \ell_t \rangle^2 \langle P_t^{-1} x_t, P_t P_t^{-1} x_t \rangle \\ &\leq \langle P_t^{-1} x_t, x_t \rangle. \end{aligned}$$

Now we use a spectral decomposition of P_t in an orthonormal basis for $\langle \cdot, \cdot \rangle$ and write $P_t = \sum_{i=1}^d \lambda_i v_i \otimes v_i$. In particular, we have $P_t^{-1} = \sum_{i=1}^d \frac{1}{\lambda_i} v_i \otimes v_i$ and thus

$$\begin{aligned} \mathbb{E} \langle P_t^{-1} x_t, x_t \rangle &= \sum_{i=1}^d \frac{1}{\lambda_i} \mathbb{E} \langle (v_i \otimes v_i) x_t, x_t \rangle \\ &= \sum_{i=1}^d \frac{1}{\lambda_i} \mathbb{E} \langle (x_t \otimes x_t) v_i, v_i \rangle \\ &= \sum_{i=1}^d \frac{1}{\lambda_i} \langle P_t v_i, v_i \rangle \\ &= \sum_{i=1}^d \frac{1}{\lambda_i} \langle \lambda_i v_i, v_i \rangle = d. \end{aligned}$$

Finally, to show $\eta \langle x, \tilde{\ell}_t \rangle \leq 1$ observe that

$$\langle x, \tilde{\ell}_t \rangle = \langle x_t, \ell_t \rangle \langle x, P_t^{-1} x_t \rangle \leq \langle x, P_t^{-1} x_t \rangle \leq \frac{1}{\min_{1 \leq i \leq d} \lambda_i},$$

where the last inequality follows from the fact that $\langle x, x \rangle \leq 1$ for any $x \in \mathcal{K}$, since \mathcal{K} is included in the unit ball. To conclude the proof, we need to lower bound the smallest eigenvalue of P_t . Using Theorem 5.1, one can see that $P_t \succeq \frac{\gamma}{d} I_d$, and thus $\lambda_i \geq \frac{\gamma}{d}$. Since $\eta d \leq \gamma$, the proof is concluded. \square

5.2 Online Mirror Descent (OMD)

We now introduce the Online Mirror Descent (OMD) algorithm, a powerful generalization of gradient descent for sequential decision problems. We start by describing OMD for convex losses in the full information setting. That is, \mathcal{L} is a set of convex functions, and at the end of round t the forecaster observes $\ell_t \in \mathcal{L}$ rather than only $\ell_t(x_t)$.

The rest of this section is organized as follows. Next, we briefly recall a few key concepts from convex analysis. Then we describe the OMD strategy and prove a general regret bound. In Section 5.3 we introduce Online Stochastic Mirror Descent (OMSD), which is a variant of OMD based on a stochastic estimate of the gradient. We apply this strategy to linear losses in two different bandit settings. Finally, in Section 5.5 we show how OMSD obtains improved bounds in certain special cases. The case of convex losses with bandit feedback is treated in Section 6.

We introduce the following definitions.

Definition 5.1. Let $\mathcal{X} \subseteq \mathbb{R}^d$. A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is subdifferentiable if for all $x \in \mathcal{X}$ there exists $g \in \mathbb{R}^d$ such that

$$f(x) - f(y) \leq g^\top(x - y) \quad \forall y \in \mathcal{X}.$$

Such a g is called a subgradient of f at x .

Abusing notation, we use $\nabla f(x)$ to denote both the gradient of f at x when f is differentiable and any subgradient of f at x when f is subdifferentiable (a sufficient condition for subdifferentiability of f is that f is convex and \mathcal{X} is open).

Definition 5.2. Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be a convex function defined on a convex set $\mathcal{X} \subseteq \mathbb{R}^d$. The Legendre–Fenchel transform of f is defined by:

$$f^*(u) = \sup_{x \in \mathcal{X}} (x^\top u - f(x)).$$

This definition directly implies the Fenchel–Young inequality for convex functions, $u^\top x \leq f(x) + f^*(u)$.

Let $\mathcal{D} \subset \mathbb{R}^d$ be an open convex set, and let $\overline{\mathcal{D}}$ be the closure of \mathcal{D} .

Definition 5.3. A continuous function $F : \overline{\mathcal{D}} \rightarrow \mathbb{R}$ is Legendre if

- (i) F is strictly convex and admits continuous first partial derivatives on \mathcal{D} ;
 - (ii) $\lim_{x \rightarrow \overline{\mathcal{D}} \setminus \mathcal{D}} \|\nabla F(x)\| = +\infty$.¹
-

The Bregman divergence $D_F : \overline{\mathcal{D}} \times \mathcal{D} \rightarrow \mathbb{R}$ associated with a Legendre function F is defined by $D_F(x, y) = F(x) - F(y) - (x - y)^\top \nabla F(y)$. Moreover, we say that $\mathcal{D}^* = \nabla F(\mathcal{D})$ is the dual space of \mathcal{D} under F . Note that, by definition, $D_F(x, y) > 0$ if $x \neq y$, and $D_F(x, x) = 0$. The following lemma is the key to understand how a Legendre function acts on the space. See [61, Proposition 11.1] for a proof.

Lemma 5.3. Let F be a Legendre function. Then $F^{**} = F$, and $\nabla F^* = (\nabla F)^{-1}$ restricted on the set \mathcal{D}^* . Moreover, for all $x, y \in \mathcal{D}$,

$$D_F(x, y) = D_{F^*}(\nabla F(y), \nabla F(x)). \quad (5.2)$$

The gradient ∇F maps \mathcal{D} to the dual space \mathcal{D}^* , and ∇F^* is the inverse mapping from the dual space to the original (primal) space. Note also that Equation (5.2) shows that the Bregman divergence in the primal exactly corresponds to the Bregman divergence of the Legendre transform in the dual.

The next lemma shows that the geometry induced by a Bregman divergence resembles to the geometry of the squared Euclidean distance. See [61, Lemma 11.3] for a proof.

Lemma 5.4 (Generalized Pythagorean inequality). Let $\mathcal{K} \subseteq \overline{\mathcal{D}}$ be a closed convex set such that $\mathcal{K} \cap \mathcal{D} \neq \emptyset$. Then, for all $x \in \mathcal{D}$, the Bregman projection

$$z = \operatorname{argmin}_{y \in \mathcal{K}} D_F(y, x)$$

¹By the equivalence of norms in \mathbb{R}^d , this definition does not depend on the choice of the norm.

exists and is unique. Moreover, for all $z \in \mathcal{K} \cap \mathcal{D}$ and $y \in \mathcal{K}$,

$$D_F(y, x) \geq D_F(y, z) + D_F(z, x).$$

The idea of OMD is very simple: first, select a Legendre function F on $\overline{\mathcal{D}} \supset \mathcal{K}$ (such that $\mathcal{K} \cap \mathcal{D} \neq \emptyset$); second, perform a gradient descent step, where the update with the gradient is performed in the dual space \mathcal{D}^* rather than in the primal \mathcal{D} ; and third, project back to \mathcal{K} according to the Bregman divergence defined by F .

OMD (Online Mirror Descent):

Parameters: compact and convex set $\mathcal{K} \subseteq \mathbb{R}^d$, learning rate $\eta > 0$, Legendre function F on $\overline{\mathcal{D}} \supset \mathcal{K}$.

Initialize: $x_1 \in \operatorname{argmin}_{x \in \mathcal{K}} F(x)$ (note that $x_1 \in \mathcal{K} \cap \mathcal{D}$).

For each round $t = 1, 2, \dots, n$

- (1) Play x_t and observe loss vector ℓ_t .
- (2) $w_{t+1} = \nabla F^*(\nabla F(x_t) - \eta \nabla \ell_t(x_t))$.
- (3) $x_{t+1} = \operatorname{argmin}_{y \in \mathcal{K}} D_F(y, w_{t+1})$.

Note that step (2) is well defined if the following consistency condition is satisfied:

$$\nabla F(x) - \eta \nabla \ell(x) \in \mathcal{D}^* \quad \forall (x, \ell) \in (\mathcal{K} \cap \mathcal{D}) \times \mathcal{L}. \quad (5.3)$$

Note also that step (2) can be rewritten as

$$\nabla F(w_{t+1}) = \nabla F(x_t) - \eta \nabla \ell_t(x_t). \quad (5.4)$$

Finally, note that the Bregman projection in step (3) is a convex program, in the sense that $x \mapsto D_F(x, y)$ is always a convex function. This does not necessarily imply that step (3) can be performed efficiently, since in some cases the feasible set \mathcal{K} might only be described with an exponential number of constraints (we encounter examples like this in Section 5.4).

In the description above we emphasized that F has to be a Legendre function. In fact, as we see in Theorem 5.6, if F has effective domain

\mathcal{K} (that is, F takes value $+\infty$ outside of \mathcal{K}), then it suffices that the Legendre–Fenchel dual F^* is differentiable on \mathbb{R}^d to obtain a good regret bound. See the end of this section for more details on this.

When \mathcal{K} is the simplex and $F(x) = \sum_{i=1}^d x_i \ln x_i - \sum_{i=1}^d x_i$, OMD reduces to an exponentially weighted average forecaster, similar to those studied in Section 3. The well-known online gradient descent strategy corresponds to taking $F(x) = \frac{1}{2} \|x\|_2^2$. In the following we shall see other possibilities for the Legendre function F .

We prove now a very general and powerful theorem concerning the regret of OMD.

Theorem 5.5 (Regret of OMD with a Legendre function). Let \mathcal{K} be a compact and convex set of arms, \mathcal{L} be a set of subdifferentiable functions, and F a Legendre function defined on $\bar{\mathcal{D}} \supset \mathcal{K}$, such that Equation (5.3) is satisfied. Then OMD satisfies for any $x \in \mathcal{K}$,

$$\begin{aligned} \sum_{t=1}^n \ell_t(x_t) - \sum_{t=1}^n \ell_t(x) &\leq \frac{F(x) - F(x_1)}{\eta} \\ &\quad + \frac{1}{\eta} \sum_{t=1}^n D_{F^*}(\nabla F(x_t) - \eta \nabla \ell_t(x_t), \nabla F(x_t)). \end{aligned}$$

Proof. Let $x \in \mathcal{K}$. Since \mathcal{L} is a set of subdifferentiable functions, we have:

$$\sum_{t=1}^n (\ell_t(x_t) - \ell_t(x)) \leq \sum_{t=1}^n \nabla \ell_t(x_t)^\top (x_t - x).$$

Using Equation (5.4), and applying the definition of D_F , one obtains

$$\begin{aligned} \eta \nabla \ell_t(x_t)^\top (x_t - x) &= (x - x_t)^\top (\nabla F(w_{t+1}) - \nabla F(x_t)) \\ &= D_F(x, x_t) + D_F(x_t, w_{t+1}) - D_F(x, w_{t+1}). \end{aligned}$$

By Lemma 5.4, one gets $D_F(x, w_{t+1}) \geq D_F(x, x_{t+1}) + D_F(x_{t+1}, w_{t+1})$, hence

$$\begin{aligned} \eta \nabla \ell_t(x_t)^\top (x_t - x) &\leq D_F(x, x_t) + D_F(x_t, w_{t+1}) \\ &\quad - D_F(x, x_{t+1}) - D_F(x_{t+1}, w_{t+1}). \end{aligned}$$

Summing over t then gives

$$\begin{aligned} \sum_{t=1}^n \eta \nabla \ell_t(x_t)^\top (x_t - x) &\leq D_F(x, x_1) - D_F(x, x_{n+1}) \\ &\quad + \sum_{t=1}^n (D_F(x_t, w_{t+1}) - D_F(x_{t+1}, w_{t+1})). \end{aligned}$$

By the nonnegativity of the Bregman divergences, we get

$$\sum_{t=1}^n \eta \nabla \ell_t(x_t)^\top (x_t - x) \leq D_F(x, x_1) + \sum_{t=1}^n D_F(x_t, w_{t+1}).$$

From Equation (5.2), one has $D_F(x_t, w_{t+1}) = D_{F^*}(\nabla F(x_t) - \eta \nabla \ell_t(x_t), \nabla F(a_t))$ and, moreover, by first-order optimality, one has $D_F(x, x_1) \leq F(x) - F(x_1)$, which concludes the proof. \square

We show now how to prove a regret bound if F has effective domain \mathcal{K} and F^* is differentiable on \mathbb{R}^d , but not necessarily Legendre. In this case, it is easy to see that step (1) and step (2) in OMD reduce to

$$x_{t+1} = \nabla F^* \left(-\eta \sum_{s=1}^{t-1} \nabla \ell_s(x_s) \right).$$

Theorem 5.6 (Regret of OMD with non-Legendre function).

Let \mathcal{K} be a compact set of actions, \mathcal{L} be a set of subdifferentiable functions, and F a function with effective domain \mathcal{K} such that F^* is differentiable on \mathbb{R}^d . Then for any $x \in \mathcal{K}$ OMD satisfies

$$\begin{aligned} \sum_{t=1}^n \ell_t(x_t) - \sum_{t=1}^n \ell_t(x) &\leq \frac{F(x) - F(x_1)}{\eta} \\ &\quad + \frac{1}{\eta} \sum_{t=1}^n D_{F^*} \left(-\eta \sum_{s=1}^t \nabla \ell_s(x_s), -\eta \sum_{s=1}^{t-1} \nabla \ell_s(x_s) \right). \end{aligned}$$

Proof. Let $x \in \mathcal{K}$. Since \mathcal{L} is a set of subdifferentiable functions, we have

$$\sum_{t=1}^n (\ell_t(x_t) - \ell_t(x)) \leq \sum_{t=1}^n \nabla \ell_t(x_t)^\top (x_t - x).$$

Using the Fenchel–Young inequality, one obtains

$$\begin{aligned}
& -\eta \sum_{t=1}^n \nabla \ell_t(x_t)^\top x \\
& \leq F(x) + F^* \left(-\eta \sum_{t=1}^n \nabla \ell_t(x_t) \right) \\
& = F(a) + F^*(0) \\
& \quad + \sum_{t=1}^n \left(F^* \left(-\eta \sum_{s=1}^t \nabla \ell_s(x_s) \right) - F^* \left(-\eta \sum_{s=1}^{t-1} \nabla \ell_s(x_s) \right) \right).
\end{aligned}$$

Observe that $F^*(0) = -F(x_1)$ and, for each term in the above sum,

$$\begin{aligned}
& \nabla F^* \left(-\eta \sum_{s=1}^{t-1} \nabla \ell_s(x_s) \right)^\top (-\eta \nabla \ell_t(x_t)) \\
& \quad + D_{F^*} \left(-\eta \sum_{s=1}^t \nabla \ell_s(x_s), -\eta \sum_{s=1}^{t-1} \nabla \ell_s(x_s) \right) \\
& = -\eta x_t^\top \nabla \ell_t(x_t) + D_{F^*} \left(-\eta \sum_{s=1}^t \nabla \ell_s(x_s), -\eta \sum_{s=1}^{t-1} \nabla \ell_s(x_s) \right). \quad \square
\end{aligned}$$

5.3 Online Stochastic Mirror Descent (OSMD)

We now turn to the analysis of the bandit setting, where the gradient information $\nabla \ell_t(x_t)$ is not available, and thus one cannot run OMD. This scenario has been extensively in gradient-free optimization, and the basic idea is to play a perturbed version \tilde{x}_t of the current point x_t . This should be done in such a way that, upon observing $\ell_t(\tilde{x}_t)$, one can build an unbiased estimate \tilde{g}_t of $\nabla \ell_t(x_t)$. Whereas such estimators are presented in Section 6, here we restrict our attention to linear losses. For this simpler case we consider specialized estimators with optimal performances. Given a perturbation scheme, one can run OMD with the gradient estimates instead of the true gradients. We call Online Stochastic Mirror Descent (OSMD) the resulting algorithm.

OSMD (Online Stochastic Mirror Descent):

Parameters: compact and convex set $\mathcal{K} \subseteq \mathbb{R}^d$, learning rate $\eta > 0$, Legendre function F on $\overline{\mathcal{D}} \supset \mathcal{K}$.

Initialize: $x_1 \in \operatorname{argmin}_{x \in \mathcal{K}} F(x)$ (note that $x_1 \in \mathcal{K} \cap \mathcal{D}$).

For each round $t = 1, 2, \dots, n$

- (1) Play a random perturbation \tilde{x}_t of x_t and observe $\ell_t(\tilde{x}_t)$
- (2) Compute random estimate \tilde{g}_t of $\nabla \ell_t(x_t)$
- (3) $w_{t+1} = \nabla F^*(\nabla F(x_t) - \eta \tilde{g}_t)$
- (4) $x_{t+1} = \operatorname{argmin}_{y \in \mathcal{K}} D_F(y, w_{t+1})$

In order to relate this linear bandit strategy to the Exp2 forecaster (Equation (5.1)), it is important to observe that running the Exp2 forecaster over a finite set \mathcal{K} of arms, with exploration distribution μ and mixing coefficient $\gamma > 0$, is equivalent to running OSMD over the $|\mathcal{K}|$ -dimensional simplex with $F(x) = \frac{1}{\eta} \sum_{x \in \mathcal{K}} x \ln x$ (the negative entropy), \tilde{x}_t drawn from $(1 - \gamma)x_t + \gamma\mu$, and estimated linear loss $\tilde{g}_t = (\langle x, \tilde{\ell}_t \rangle)_{x \in \mathcal{K}}$. Indeed, the projection step (4), when F is the negative entropy, corresponds to the standard normalization of a probability distribution.

The following theorem establishes a general regret bound for OSMD. Note that here the pseudo-regret is defined as

$$\overline{R}_n = \mathbb{E} \sum_{t=1}^n \ell_t(\tilde{x}_t) - \min_{x \in \mathcal{K}} \mathbb{E} \sum_{t=1}^n \ell_t(x).$$

Note also that we state the theorem for a Legendre function F , but a similar result can be obtained under the same assumptions as those of Theorem 5.6.

Theorem 5.7 (Pseudo-regret of OSMD). Let \mathcal{K} be a compact and convex set, \mathcal{L} a set of subdifferentiable functions, and F a Legendre function defined on $\overline{\mathcal{D}} \supset \mathcal{K}$. If OSMD is run with a loss estimate \tilde{g}_t such that Equation (5.3) is satisfied (with $\nabla \ell(x)$ replaced by \tilde{g}_t), and

with $\mathbb{E}[\tilde{g}_t | x_t] = \nabla \ell_t(x_t)$, then

$$\begin{aligned} \bar{R}_n &\leq \frac{\sup_{x \in \mathcal{K}} F(x) - F(x_1)}{\eta} + \frac{1}{\eta} \sum_{t=1}^n \mathbb{E} D_{F^*}(\nabla F(x_t) - \eta \tilde{g}_t, \nabla F(x_t)) \\ &\quad + \sum_{t=1}^n \mathbb{E}[\|x_t - \tilde{x}_t\| \|\tilde{g}_t\|_*] \end{aligned}$$

for any norm $\|\cdot\|$. Moreover if the loss is linear, that is $\ell(x) = \ell^\top x$, then

$$\begin{aligned} \bar{R}_n &\leq \frac{\sup_{x \in \mathcal{K}} F(x) - F(x_1)}{\eta} + \frac{1}{\eta} \sum_{t=1}^n \mathbb{E} D_{F^*}(\nabla F(x_t) - \eta \tilde{g}_t, \nabla F(x_t)) \\ &\quad + \sum_{t=1}^n \mathbb{E}[\|x_t - \mathbb{E}[\tilde{x}_t | x_t]\| \|\tilde{g}_t\|_*]. \end{aligned}$$

Proof. Using Theorem 5.5 one directly obtains:

$$\sum_{t=1}^n \tilde{g}_t^\top (x_t - x) \leq \frac{F(x) - F(x_1)}{\eta} + \frac{1}{\eta} \sum_{t=1}^n \mathbb{E} D_{F^*}(\nabla F(x_t) - \eta \tilde{g}_t, \nabla F(x_t)).$$

Moreover since $\mathbb{E}[\tilde{g}_t | x_t] = \nabla \ell_t(x_t)$, one has:

$$\begin{aligned} \mathbb{E} \sum_{t=1}^n (\ell_t(\tilde{x}_t) - \ell_t(x)) &= \mathbb{E} \sum_{t=1}^n (\ell_t(\tilde{x}_t) - \ell_t(x_t) + \ell_t(x_t) - \ell_t(x)) \\ &\leq \mathbb{E} \sum_{t=1}^n \|x_t - \tilde{x}_t\| \|\tilde{g}_t\|_* + \mathbb{E} \sum_{t=1}^n \nabla \ell_t(x_t)^\top (x_t - x) \\ &= \mathbb{E} \sum_{t=1}^n \|x_t - \tilde{x}_t\| \|\tilde{g}_t\|_* + \mathbb{E} \sum_{t=1}^n \tilde{g}_t^\top (x_t - x) \end{aligned}$$

which concludes the proof of the first regret bound. The case of a linear loss follows very easily from the same computations. \square

5.4 Online Combinatorial Optimization

In this section we consider an interesting special case of online linear optimization. In the online combinatorial optimization setting the set of arms is $\mathcal{C} \subseteq \{0, 1\}^d$ and the set of linear loss functions is $\mathcal{L} = [0, 1]^d$. We

assume $\|v\|_1 = m$ for all $v \in \mathcal{C}$ and for some integer $m \leq d$. Many interesting problems fall into this framework, including ranking/selection of m items, or path planning.

Here we focus on the version of the problem with *semi-bandit* feedback, which is defined as follows: after playing $v_t \in \mathcal{C}$, one observes $(\ell_t(1)v_t(1), \dots, \ell_t(d)v_t(d))$. Namely, one only observes the coordinates of the loss that were *active* in the arm v_t that we chose. This setting has thus a much weaker feedback than the full information case, but still stronger than the bandit case. Note that the semi-bandit setting includes the basic multi-armed bandit problem of Section 3, which simply corresponds to $\mathcal{C} = \{e_1, \dots, e_d\}$ where e_1, \dots, e_d is the canonical basis of \mathbb{R}^d .

Again, the key to tackle this kind of problem is to select v_t at random from some probability distribution p_t over \mathcal{C} . Note that such a probability corresponds to an average point $x_t \in \text{Conv}(\mathcal{C})$. Turning the tables, one can view v_t as a random perturbation of x_t such that $\mathbb{E}[v_t | x_t] = x_t$. This suggests a strategy: play OSMD on $\mathcal{K} = \text{Conv}(\mathcal{C})$, with $\tilde{x}_t = v_t$. Surprisingly, we show that this randomization is enough to obtain a good unbiased estimate of the loss, and that it is not necessary to add further perturbations to x_t . Note that $\mathbb{E}[\tilde{x}_t | x_t] = x_t$ by definition. We now need to describe how to obtain an unbiased estimate of the gradient (which is the loss vector itself, since losses are linear). The following simple formula gives an unbiased estimate of the loss:

$$\tilde{\ell}_t(i) = \frac{\ell_t(i)v_t(i)}{x_t(i)} \quad \forall i \in \{1, \dots, d\}. \quad (5.5)$$

Note that this is a valid estimate since it only makes the use of $(\ell_t(1)v_t(1), \dots, \ell_t(d)v_t(d))$. Moreover, it is unbiased with respect to the random drawing of v_t from p_t . Indeed,

$$\mathbb{E}[\tilde{\ell}_t(i) | x_t] = \frac{\ell_t(i)}{x_t(i)} \mathbb{E}[v_t(i) | x_t] = \ell_t(i).$$

Using Theorem 5.7 one directly obtains:

$$\begin{aligned} \bar{R}_n &\leq \frac{\sup_{x \in \mathcal{K}} F(x) - F(x_1)}{\eta} \\ &\quad + \frac{1}{\eta} \sum_{t=1}^n \mathbb{E} D_{F^*}(\nabla F(x_t) - \eta \tilde{\ell}_t, \nabla F(x_t)). \end{aligned} \quad (5.6)$$

We show now how to use this bound to obtain concrete performances for OSMD using the negative entropy as Legendre function. Later, we show that one can improve the results by logarithmic factors, using a more subtle Legendre function.

We start with OSMD and the negative entropy.

Theorem 5.8 (OSMD with negative entropy). For any set $\mathcal{C} \subseteq \{0, 1\}^d$, if OSMD is run on $\mathcal{K} = \text{Conv}(\mathcal{C})$ with $F(x) = \sum_{i=1}^d x_i \ln x_i - \sum_{i=1}^d x_i$, perturbed points \tilde{x}_t such that $\mathbb{E}[\tilde{x}_t | x_t] = x_t$, and loss estimates $\tilde{\ell}_t$, then

$$\bar{R}_n \leq \frac{m}{\eta} \ln \frac{d}{m} + \frac{\eta}{2} \sum_{t=1}^n \sum_{i=1}^d \mathbb{E}[x_t(i) \tilde{\ell}_t(i)^2].$$

In particular, with the estimate (Equation (5.5)) and $\eta = \sqrt{\frac{2m}{nd} \ln \frac{d}{m}}$,

$$\bar{R}_n \leq \sqrt{2mdn \ln \frac{d}{m}}.$$

Proof. First note that:

$$F(x) - F(x_1) \leq \sum_{i=1}^d x_1(i) \ln \frac{1}{x_1(i)} \leq m \ln \left(\sum_{i=1}^d \frac{x_1(i)}{m} \frac{1}{x_1(i)} \right) = m \ln \frac{d}{m}.$$

Moreover, straightforward computations give

$$D_{F^*}(\nabla F(x_t) - \eta \tilde{\ell}_t, \nabla F(x_t)) = \sum_{i=1}^d x_t(i) \Theta(-\eta \tilde{\ell}_t(i))$$

where $\Theta : x \in \mathbb{R} \mapsto \exp(x) - 1 - x$. Using that $\Theta(x) \leq \frac{x^2}{2}$ for all $x \leq 0$, concludes the proof of the first inequality (since $\tilde{\ell}_t(i) \geq 0$). The second inequality follows from

$$x_t(i) \mathbb{E}[\tilde{\ell}_t(i)^2 | x_t] = x_t(i) \frac{\ell_t(i)^2}{x_t(i)^2} \mathbb{E}[v_t(i) | x_t] \leq 1$$

where we used $\ell_t(i) \in [0, 1]$ and $v_t(i) \in \{0, 1\}$. □

We now greatly generalize the negative entropy with the following definition. When used with OSMD, this more general entropy allows us to obtain a bound tighter than that of Theorem 5.8.

Definition 5.4. Let $\omega \geq 0$. A function $\psi : (-\infty, a) \rightarrow \mathbb{R}_+^*$ for some $a \in \mathbb{R} \cup \{+\infty\}$ is called an ω -potential if it is convex, continuously differentiable, and satisfies

$$\begin{aligned} \lim_{x \rightarrow -\infty} \psi(x) = \omega & & \lim_{x \rightarrow a} \psi(x) = +\infty \\ \psi' > 0 & & \int_{\omega}^{\omega+1} |\psi^{-1}(s)| ds < +\infty. \end{aligned}$$

With a potential ψ we associate the function F_ψ defined on $\mathcal{D} = (\omega, +\infty)^d$ by

$$F_\psi(x) = \sum_{i=1}^d \int_{\omega}^{x_i} \psi^{-1}(s) ds.$$

We restrict our attention to 0-potentials. A nonzero ω might be used to derive high probability regret bounds (instead of pseudo-regret bounds). Note that with $\psi(x) = e^x$ we have that F_ψ reduces to the negative entropy.

Lemma 5.9. Let ψ be a 0-potential. Then F_ψ is Legendre and for all $u, v \in \mathcal{D}^* = (-\infty, a)^d$ such that $u_i \leq v_i$ for $i = 1, \dots, d$,

$$D_{F^*}(u, v) \leq \frac{1}{2} \sum_{i=1}^d \psi'(v_i) (u_i - v_i)^2.$$

Proof. It is easy to check that F is a Legendre function. Moreover, since $\nabla F^*(u) = (\nabla F)^{-1}(u) = (\psi(u_1), \dots, \psi(u_d))$ we obtain

$$D_{F^*}(u, v) = \sum_{i=1}^d \left(\int_{v_i}^{u_i} \psi(s) ds - (u_i - v_i) \psi(v_i) \right).$$

From a Taylor expansion, we have

$$D_{F^*}(u, v) \leq \sum_{i=1}^d \max_{s \in [u_i, v_i]} \frac{1}{2} \psi'(s) (u_i - v_i)^2.$$

Since the function ψ is convex, and $u_i \leq v_i$, we have

$$\max_{s \in [u_i, v_i]} \psi'(s) \leq \psi'(\max\{u_i, v_i\}) \leq \psi'(v_i)$$

which gives the desired result. \square

We are now ready to bound the pseudo-regret of OSMD run with an arbitrary 0-potential. For a specific choice of the potential we obtain an improvement of Theorem 5.8. In particular for $m = 1$ this result gives the log-free bound for the adversarial multi-armed bandit that was discussed in Section 3.4.1.

Theorem 5.10 (OSMD with a 0-potential). For any set subset \mathcal{C} of $\{0, 1\}^d$, if OSMD is run on $\mathcal{K} = \text{Conv}(\mathcal{C})$ with F_ψ defined by a 0-potential ψ , and nonnegative loss estimates $\tilde{\ell}_t$, then

$$\bar{R}_n \leq \frac{\sup_{x \in \mathcal{K}} F_\psi(x) - F_\psi(x_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^n \sum_{i=1}^d \mathbb{E} \left[\frac{\tilde{\ell}_t(i)^2}{(\psi^{-1})'(x_t(i))} \right].$$

In particular, choosing the 0-potential $\psi(x) = (-x)^{-q}$, the estimate (Equation (5.5)), and $\eta = \sqrt{\frac{2}{q-1} \frac{m^{1-2/q}}{d^{1-2/q}}}$,

$$\bar{R}_n \leq q \sqrt{\frac{2}{q-1} m d n}.$$

With $q = 2$ this gives

$$\bar{R}_n \leq 2\sqrt{2m d n}.$$

Proof. First note that since $\mathcal{D}^* = (-\infty, a)^d$ and $\tilde{\ell}_t$ has nonnegative coordinates, then Equation (5.3) is satisfied and thus OSMD is well defined.

The first inequality trivially follows from Equation (5.6), Lemma 5.9, and the fact that $\psi'(\psi^{-1}(s)) = \frac{1}{(\psi^{-1})'(s)}$.

Let $\psi(x) = (-x)^{-q}$. Then we have that $\psi^{-1}(x) = -x^{-1/q}$ and $F(x) = -\frac{q}{q-1} \sum_{i=1}^d x_i^{1-1/q}$. In particular, by Hölder's inequality, since $\sum_{i=1}^d x_1(i) = m$,

$$F(x) - F(x_1) \leq \frac{q}{q-1} \sum_{i=1}^d x_1(i)^{1-1/q} \leq \frac{q}{q-1} m^{(q-1)/q} d^{1/q}.$$

Moreover, note that $(\psi^{-1})'(x) = \frac{1}{q} x^{-1-1/q}$, and

$$\sum_{i=1}^d \mathbb{E} \left[\frac{\tilde{\ell}_t(i)^2}{(\psi^{-1})'(x_t(i))} \middle| x_t \right] \leq q \sum_{i=1}^d x_t(i)^{1/q} \leq qm^{1/q} d^{1-1/q}$$

which ends the proof. \square

5.5 Improved Regret Bounds for Bandit Feedback

We go back to the setting of linear losses with bandit feedback considered in Section 5.1. Namely, actions belong to a compact and convex set $\mathcal{K} \subseteq \mathbb{R}^d$, losses belong to a subset $\mathcal{L} \subseteq \mathbb{R}^d$, and the loss of playing $x_t \in \mathcal{K}$ at time t is $x_t^\top \ell_t$, which is also the feedback received by the player. As we proved in Section 5.1, under the bounded scalar loss assumption, $|x^\top \ell| \leq 1$ for all $(x, \ell) \in \mathcal{K} \times \mathcal{L}$, one can obtain a regret bound of order $d\sqrt{n}$ (up to logarithmic factors) for any compact and convex set \mathcal{K} . It can be shown that this rate is not improvable in general. However, results from Section 5.4 (or from Section 3) show that for the simplex, one can obtain a regret bound of order \sqrt{dn} , and we showed in Section 3 that this rate is also unimprovable. The problem of obtaining a characterization of the sets for which such improved regret bounds are possible is an open problem. Improved rates can be obtained for another convex body: the Euclidean ball. We now describe a strategy that attains a pseudo-regret of order \sqrt{dn} (up to a logarithmic factor) in this case. The strategy is based on OSMD with a carefully chosen Legendre function.

In the following, let $\|\cdot\|$ be the Euclidean norm. We consider the online linear optimization problem with bandit feedback on the

Euclidean unit ball $\mathcal{K} = \{x \in \mathbb{R}^d : \|x\| \leq 1\}$. We perform the following perturbation of a point x_t in the interior of \mathcal{K} ,

$$\tilde{x}_t = \begin{cases} x_t / \|x_t\| & \text{if } \xi_t = 1, \\ \varepsilon_t e_{I_t} & \text{otherwise} \end{cases}$$

where ξ_t is a Bernoulli random variable of parameter $\|x_t\|$, I_t is drawn uniformly at random in $\{1, \dots, d\}$, and ε_t is a Rademacher random variable with parameter $\frac{1}{2}$.

It is easy to check that this perturbation is unbiased, in the sense that $\mathbb{E}[\tilde{x}_t | x_t] = x_t$. An unbiased estimate of the loss vector is given by

$$\tilde{\ell}_t = d(1 - \xi_t) \frac{\tilde{x}_t^\top \ell_t}{1 - \|x_t\|} \tilde{x}_t. \quad (5.7)$$

Again, it is easy to check that $\mathbb{E}[\tilde{\ell}_t | x_t] = x_t$. We are now ready to prove the following result, showing that OSMD with a suitable F achieves a pseudo-regret of order $\sqrt{dn \ln n}$ on the Euclidean ball.

Theorem 5.11 (OSMD for the Euclidean ball). Let $\mathcal{K} = \mathcal{L} = \{x \in \mathbb{R}^d : \|x\| \leq 1\}$ define an online linear optimization problem with bandit feedback. If OSMD is run on $\mathcal{K}' = (1 - \gamma)\mathcal{K}$ with $F(x) = -\ln(1 - \|x\|) - \|x\|$ and the estimate (Equation (5.7)), then for any $\eta > 0$ such that $\eta d \leq \frac{1}{2}$,

$$\bar{R}_n \leq \gamma n + \frac{\ln \gamma^{-1}}{\eta} + \eta \sum_{t=1}^n \mathbb{E}[(1 - \|x_t\|) \|\tilde{\ell}_t\|^2]. \quad (5.8)$$

In particular, with $\gamma = \frac{1}{\sqrt{n}}$ and $\eta = \sqrt{\frac{\ln n}{2nd}}$,

$$\bar{R}_n \leq 3\sqrt{dn \ln n}. \quad (5.9)$$

Proof. First, it is clear that by playing on $\mathcal{K}' = (1 - \gamma)\mathcal{K}$ instead of \mathcal{K} , OSMD incurs an extra γn regret. Second, note that F is strictly convex (it is the composition of a convex and nondecreasing function with the Euclidean norm) and

$$\nabla F(x) = \frac{x}{1 - \|x\|}. \quad (5.10)$$

In particular, F is Legendre on the open unit ball $\mathcal{D} = \{x \in \mathbb{R}^d : \|x\| < 1\}$, and one has $\mathcal{D}^* = \mathbb{R}^d$. Hence Equation (5.3) is always satisfied and OSMD is well defined. Now the regret with respect to \mathcal{K}' can be bounded as follows: using Theorem 5.7 and the unbiasedness of \tilde{x}_t and $\tilde{\ell}_t$ we get

$$\frac{\sup_{x \in \mathcal{K}} F(x) - F(x_1)}{\eta} + \frac{1}{\eta} \sum_{t=1}^n \mathbb{E} D_{F^*}(\nabla F(x_t) - \eta \tilde{\ell}_t, \nabla F(x_t)).$$

The first term is clearly bounded by $\frac{1}{\eta} \ln \frac{1}{\gamma}$ (since $x_1 = 0$). For the second term, we need to do a few computations. The first one follows from Equation (5.10), the others follow from simple algebra

$$\begin{aligned} \nabla F^*(u) &= \frac{u}{1 + \|u\|} \\ F^*(u) &= -\ln(1 + \|u\|) + \|u\| \\ D_{F^*}(u, v) &= \frac{1}{1 + \|v\|} \left(\|u\| - \|v\| + \|u\| \cdot \|v\| - v^\top u \right. \\ &\quad \left. - (1 + \|v\|) \ln \left(1 + \frac{\|u\| - \|v\|}{1 + \|v\|} \right) \right). \end{aligned}$$

Let $\Theta(u, v)$ such that $D_{F^*}(u, v) = \frac{1}{1 + \|v\|} \Theta(u, v)$. First note that

$$\frac{1}{1 + \|\nabla F(x_t)\|} = 1 - \|x_t\|. \quad (5.11)$$

Thus, in order to prove Equation (5.8) it remains to show that $\Theta(u, v) \leq \|u - v\|^2$, for $u = \nabla F(x_t) - \eta \tilde{\ell}_t$ and $v = \nabla F(x_t)$. In fact, we prove that this inequality holds as soon as $\frac{\|u\| - \|v\|}{1 + \|v\|} \geq -\frac{1}{2}$. This is the case for the pair (u, v) under consideration, since by the triangle inequality, Equations (5.7) and (5.11), and the assumption on η ,

$$\frac{\|u\| - \|v\|}{1 + \|v\|} \geq -\frac{\eta \|\tilde{\ell}_t\|}{1 + \|v\|} \geq -\eta d \geq -\frac{1}{2}.$$

Now using that $\ln(1 + x) \geq x - x^2$ for all $x \geq -\frac{1}{2}$, we obtain that for u, v such that $\frac{\|u\| - \|v\|}{1 + \|v\|} \geq -\frac{1}{2}$,

$$\begin{aligned} \Theta(u, v) &\leq \frac{(\|u\| - \|v\|)^2}{1 + \|v\|} + \|u\| \cdot \|v\| - v^\top u \\ &\leq (\|u\| - \|v\|)^2 + \|u\| \cdot \|v\| - v^\top u \end{aligned}$$

$$\begin{aligned}
&= \|u\|^2 + \|v\|^2 - \|u\| \cdot \|v\| - v^\top u \\
&= \|u - v\|^2 + 2v^\top u - \|u\| \cdot \|v\| - v^\top u \\
&\leq \|u - v\|^2
\end{aligned}$$

which concludes the proof of Equation (5.8). For the proof of Equation (5.9) it suffices to note that

$$\begin{aligned}
\mathbb{E}[1 - \|x_t\| \|\tilde{\ell}_t\|^2] &= (1 - \|x_t\|) \sum_{i=1}^d \frac{1 - \|x_t\|}{d} \frac{d^2}{(1 - \|x_t\|)^2} (\ell_t^\top e_i)^2 \\
&= d \|\ell_t\|^2 \\
&\leq d
\end{aligned}$$

and perform with straightforward computations. \square

5.6 Refinements and Bibliographic Remarks

Online convex optimization in the full information setting was introduced by Zinkevich [175]. Online linear optimization with bandit feedback was pioneered in [27, 135]. For this problem, Dani et al. [71] were the first to obtain optimal $\mathcal{O}(\sqrt{n})$ bounds in terms of the number n of rounds. This was done using the Exp2 strategy with an exploration uniform over a barycentric spanner for \mathcal{K} . The exploration part was first improved by Cesa-Bianchi and Lugosi [62] for combinatorial sets \mathcal{K} . Finally, the optimal exploration based on John's theorem was introduced by Bubeck et al. [45]. Theorem 5.2 is extracted from this last paper.

Simultaneously with the line of research on Exp2, algorithms based on Online Mirror Descent were also investigated. Mirror Descent was originally introduced in the seminal work of Nemirovski [137] and Nemirovski and Yudin [139] as a standard (offline) convex optimization method. A somewhat similar class of algorithms was rediscovered in the online learning community, see [88, 101, 119, 151]. The connection between existing online learning algorithms (such as Exponential weights or Online Gradient Descent) and Mirror Descent was first made explicit in [61] — see also [145] and [95]. Earlier applications of Mirror Descent in the learning community can be found in [106]. The first application of Mirror Descent to online linear optimization with bandit

feedback was given by Abernethy et al. [3]. In this pioneering paper, the authors describe the first computationally efficient strategy (i.e., with complexity polynomial in d) with $\mathcal{O}(\sqrt{n})$ regret. The main idea is to use Mirror Descent with a self-concordant barrier F for the set \mathcal{K} . Unfortunately, the drawback is a suboptimal dependency on d in the regret. More precisely, they obtain a $\mathcal{O}(d^2\sqrt{n})$ regret under the bounded scalar loss assumption, while Exp2 with John’s exploration attains $\mathcal{O}(d\sqrt{n})$. However, Mirror Descent can also deliver optimal regret bounds in the bandit case, as we showed in Section 5.5, which is extracted from [45].

The presentation of the Online Mirror Descent algorithm in Section 5.2 is inspired by Bubeck [44]. The definition of Legendre functions comes from [61, Section 11] — further developments on convex analysis can be found in [42, 102]. Theorem 5.5 is taken from Audibert et al. [19], but the proof technique goes back at least to Ben-Tal and Nemirovski [37]. The proof of Theorem 5.6 is adapted from Kakade et al. [107]. Section 5.3 is inspired by gradient-free optimization, a topic extensively studied since [118, 147] — see [31, 67, 138, 140] for recent accounts on this theory. Alternative views have been proposed on the Online Mirror Descent strategy. In particular, it is equivalent to a Follow The Regularized Leader, and to proximal algorithms, see Rakhlin [145]. This viewpoint was pioneered by Beck and Teboulle [36] — see also [35] for more details. Finally, a notion of universality of Online Mirror Descent in the online prediction setting was proposed by Srebro et al. [154].

The online combinatorial optimization problem studied in Section 5.4 was introduced by Kalai and Vempala [110] for the full information setting. Several works have studied this problem for specific sets \mathcal{C} , see in particular [62, 98, 100, 124, 159, 168, 169]. The semi-bandit feedback was studied in the series of papers [19, 91, 111, 163]. The presentation adopted in this section is based on the last paper. OSMD with negative entropy was first studied by Helmbold and Warmuth [100] for the full information setting and for a specific set \mathcal{C} . It was then studied more generally in [124] for any set \mathcal{C} . The generalization to semi-bandit feedback was done by Audibert et al. [19]. OSMD with a Legendre derived from a potential was introduced by Audibert et al. [19]. In the case of the simplex, this strategy corresponds to the INF strategy of Audibert and Bubeck [17] discussed in Section 3.4.1.

Online linear optimization is still far from being completely understood. For instance, see [44, Section 9] for a list of open problems. In this section we also omitted a few important topics related to online linear optimization. We briefly review some of them below.

5.6.1 Lower Bounds

Under the bounded scalar loss assumption, it was proved by Dani et al. [71] that for $\mathcal{K} = [-1, 1]^d$ the minimax regret in the full information setting is at least of order \sqrt{dn} , while under bandit feedback it is of order $d\sqrt{n}$. In both cases Exp2 is matching these lower bounds (using John's exploration in the bandit case).

In the combinatorial setting, where $\mathcal{K} \subset \{0, 1\}^d$ and $\mathcal{L} = [0, 1]^d$, Audibert et al. [19] show that the minimax regret in the full information and semi-bandit cases is at least of order $d\sqrt{n}$, while in the bandit case it is of order $d^{3/2}\sqrt{n}$. OSMD with the negative entropy matches the bounds in the full information and semi-bandit cases. However, in the bandit case the best-known bound is obtained by Exp2 (with John's exploration) and gives a regret of order $d^2\sqrt{n}$. It is important to remark that Audibert et al. [19] show that Exp2 is a provably sub-optimal strategy in the combinatorial setting.

Finally, lower bounds for the full information case, and for a few specific sets \mathcal{K} of interest, were derived by Koolen et al. [124].

5.6.2 High Probability Bounds

In this section we focused on the pseudo-regret \bar{R}_n . However, just like in Section 3, a much more important and interesting statement concerns high probability bounds for the regret R_n . Partial results in this direction can be found in [33] for the Exp2 strategy, and in [4] for the OSMD algorithm.

5.6.3 Stochastic Online Linear Optimization

Similarly to the stochastic bandit case (see Section 2), a natural restriction to consider for the adversary is that the sequence of losses ℓ_1, ℓ_2, \dots

is an i.i.d. sequence. This stochastic setting was introduced by Auer [22], and further studied by Dani et al. [72]. In particular, in the latter paper it was proved that regrets logarithmic in n and polynomial in d are possible, as long as \mathcal{K} is a polytope. Recent progress on this problem can be found in [1, 78, 148].

6

Nonlinear Bandits

We now extend the analysis of the previous section to the following scenario: arms are still points in a convex set $\mathcal{K} \subset \mathbb{R}^d$, but now losses are not necessarily linear functions of the arms. More precisely the adversary selects loss functions from some set \mathcal{L} of real-valued functions defined on \mathcal{K} . The pseudo-regret is then defined as:

$$\bar{R}_n = \mathbb{E} \sum_{t=1}^n \ell_t(x_t) - \min_{x \in \mathcal{K}} \mathbb{E} \sum_{t=1}^n \ell_t(x).$$

This modification has important consequences. For instance with strictly convex losses one has to do local perturbations in order to estimate the loss gradient, this is in contrast to the global perturbations studied in the previous section. In agreement with the setting of Section 5, we initially focus on the nonstochastic setting, where the forecaster faces an unknown sequence of convex Lipschitz and differentiable losses (in the nonlinear case the regret scales with the Lipschitz constant of losses). Problems of this kind can be viewed as dynamic variants of convex optimization problems, in which the

convex function to optimize evolves over time. The bandit constraint can be simply interpreted as the impossibility of computing gradients (because, for instance, we do not have an explicit representation of the function, but it can only be accessed by querying for values at desired points).

We look at two feedback models. In the first one, at each step the forecaster evaluates the loss function at two points: the played point plus an additional point of its choice. In the second one, only the value of the loss evaluated at the played point is made available to the forecaster. We show that while the two-point model allows for a $\mathcal{O}(\sqrt{n})$ bound on pseudo-regret, in the one-point model a pseudo-regret bound of only $\mathcal{O}(n^{3/4})$ is achieved. The stochastic setting is investigated in Section 6.3 where, similarly to Section 2, we assume that each play of an arm returns a stochastic loss with fixed but unknown mean. Unlike the nonstochastic case, the mean loss function is assumed to be Lipschitz and unimodal, but not necessarily convex. For keeping things simple, the stochastic setting is studied in one-dimensional case, when arms are points in the unit interval. For this case we show a bound on the pseudo-regret of $\mathcal{O}(\sqrt{n}(\log n))$.

6.1 Two-Point Bandit Feedback

We start by analyzing the nonstochastic case in the two-point feedback model: at each time step t , the forecaster observes the value of a convex and differentiable loss function ℓ_t at the played point x_t and at an extra point x'_t of its choice. If the second point is chosen at random in a neighborhood of the first one, one can use it to compute an estimate of the gradient of ℓ_t at x_t . Hence, running OSMD on the estimated gradients we obtain a regret bound controlled by the second moments of these estimates. The algorithm we present in this section follows this intuition, although — for technical reasons — the gradient is estimated at a point which is close but distinct from the point actually played.

We focus our analysis on OSMD with Legendre function $F = \frac{1}{2} \|\cdot\|^2$, where $\|\cdot\|$ is the Euclidean norm. The resulting strategy, OSGD (Online Stochastic Gradient Descent), is sketched below here.

OSGD (Online Stochastic Gradient Descent):

Parameters: Closed and convex set $\mathcal{K} \subseteq \mathbb{R}^d$, learning rate $\eta > 0$.

Initialize: $x_1 = (0, \dots, 0)$.

For each round $t = 1, 2, \dots, n$

- (1) Observe stochastic estimate $\tilde{g}_t(x_t)$ of $\nabla \ell_t(x_t)$;
- (2) $x'_{t+1} = x_t - \eta \tilde{g}_t(x_t)$;
- (3) $x_{t+1} = \operatorname{argmin}_{y \in \mathcal{K}} \|y - x'_{t+1}\|$;

We now introduce our main technical tool: the two-point gradient estimate. The two points on which the loss value is queried at time t are denoted by X_t^+ and X_t^- . OSGD always plays one of these two points at random.

Let $\mathbb{B} = \{x \in \mathbb{R}^d : \|x\| \leq 1\}$ be the unit ball in \mathbb{R}^d and $\mathbb{S} = \{x \in \mathbb{R}^d : \|x\| = 1\}$ be the unit sphere. Fix $\delta > 0$ and introduce the notations $X_t^+ = x_t + \delta S$ and $X_t^- = x_t - \delta S$, where $x_t \in \mathcal{K}$ and S is a random variable with uniform distribution in \mathbb{S} . Then, for any convex loss ℓ_t , the two-point gradient estimate \tilde{g}_t is defined by

$$\tilde{g}_t(x_t) = \frac{d}{2\delta} (\ell_t(X_t^+) - \ell_t(X_t^-)) S. \quad (6.1)$$

In order to compute the expectation of \tilde{g}_t , first note that by symmetry

$$\mathbb{E} \tilde{g}_t(x) = \frac{d}{\delta} \mathbb{E} [\ell_t(x + \delta S) S].$$

In order to compute the expectation in the right-hand side we need the following preliminary lemma.

Lemma 6.1. For any differentiable function $\ell : \mathbb{R}^d \rightarrow \mathbb{R}$

$$\nabla \int_{\mathbb{B}} \ell(x + \delta b) db = \int_{\mathbb{S}} \ell(x + \delta s) s d\sigma(s)$$

where σ is the unnormalized spherical measure.

Proof. The proof of this result is an easy consequence of the Divergence Theorem,

$$\begin{aligned} \nabla \int_{\mathbb{B}} \ell(x + \delta b) db &= \int_{\mathbb{B}} \nabla \ell(x + \delta b) db \\ &= \int_{\mathbb{S}} \frac{1}{\delta} \ell(x + \delta s) s d\sigma(s) \\ &= \frac{1}{\delta} \int_{\mathbb{S}} \ell(x + \delta s) s d\sigma(s). \quad \square \end{aligned}$$

We are now fully equipped to compute the expectation of \tilde{g}_t .

Lemma 6.2. If B is a random variable with uniform distribution in \mathbb{B} and S is a random variable with uniform distribution in \mathbb{S} , then for all differentiable functions $\ell_t : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\frac{d}{\delta} \mathbb{E}[\ell(x + \delta S)S] = \nabla \mathbb{E} \ell(x + \delta B).$$

Proof. First consider the easy one-dimensional case. Namely, $\mathcal{K} = [a, b]$ for some reals $a < b$. Note that, in this case, S is uniform in $\{-1, +1\}$ whereas B is uniform in $[-1, +1]$. Then

$$\mathbb{E} \ell(x + \delta B) = \frac{1}{2\delta} \int_{-\delta}^{\delta} \ell(x + \varepsilon) d\varepsilon = \frac{L(x + \delta) - L(x - \delta)}{2\delta}$$

by the fundamental theorem of calculus, where L is the antiderivative of ℓ satisfying $L' = \ell$. This gives

$$\frac{d}{dx} \mathbb{E} \ell(x + \delta B) = \frac{\ell(x + \delta) - \ell(x - \delta)}{2\delta}.$$

On the other hand,

$$\frac{1}{\delta} \mathbb{E}[\ell(x + \delta S)S] = \frac{\ell(x + \delta) - \ell(x - \delta)}{2\delta}.$$

Hence $\frac{1}{\delta} \mathbb{E}[\ell(x + \delta S)S] = \frac{d}{dx} \mathbb{E} \ell(x + \delta B)$ and the one-dimensional case is established. Note that the equivalence we just proved relates an integral over the unit sphere \mathbb{S} to an integral over the unit ball \mathbb{B} . In d dimensions, Lemma 6.1 delivers the corresponding generalized identity

$$\frac{1}{\delta} \int_{\mathbb{S}} \ell(x + \delta s) s d\sigma(s) = \nabla \int_{\mathbb{B}} \ell(x + \delta b) db.$$

Now, since $\text{Vol}(\mathbb{S}) = d \text{Vol}(\mathbb{B})$ we immediately obtain

$$\frac{d}{\delta} \mathbb{E}[\ell(x + \delta S)S] = \nabla \mathbb{E} \ell(x + \delta B)$$

concluding the proof. \square

We have thus established $\mathbb{E} \tilde{g}_t(x) = \nabla \mathbb{E} \ell_t(x + \delta B)$, showing that \tilde{g}_t provides an unbiased estimate of a smoothed version $\tilde{\ell}_t(x) = \mathbb{E} \ell_t(x + \delta B)$ of the loss function ℓ_t .

We can measure how well $\tilde{\ell}_t$ approximates ℓ_t by exploiting the Lipschitz assumption,

$$\begin{aligned} |\ell_t(x) - \tilde{\ell}_t(x)| &= |\ell_t(x) - \mathbb{E} \ell_t(x + \delta B)| \\ &\leq \mathbb{E} |\ell_t(x) - \ell_t(x + \delta B)| \\ &\leq \delta G \mathbb{E} \|B\| \\ &\leq \delta G. \end{aligned} \tag{6.2}$$

The next lemma relates the regret under the losses ℓ_t to the regret under their smoothed versions $\tilde{\ell}_t$. An additional issue taken into account by the lemma is that OSGD might play a point close to the boundary of the set \mathcal{K} . In this case the perturbed point on which the gradient is estimated could potentially be outside of \mathcal{K} . In order to prevent this from happening we need to run OSGD on a shrunken set $(1 - \xi)\mathcal{K}$.

Lemma 6.3. Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a convex set such that $\mathcal{K} \subseteq R\mathbb{B}$ for some $R \geq 0$, and fix $0 \leq \xi \leq 1$. For any sequence ℓ_1, ℓ_2, \dots of G -Lipschitz differentiable and convex losses, and for any sequence $x_1, x_2, \dots \in (1 - \xi)\mathcal{K} \subseteq \mathbb{R}^d$, the following holds

$$\begin{aligned} &\frac{1}{2} \sum_{t=1}^n (\ell_t(X_t^+) + \ell_t(X_t^-)) - \sum_{t=1}^n \ell_t(x) \\ &\leq \sum_{t=1}^n \tilde{\ell}_t(x_t) - \sum_{t=1}^n \tilde{\ell}_t((1 - \xi)x) + 3\delta Gn + \xi GRn \end{aligned}$$

for all realizations of the random process $(X_t^+, X_t^-)_{t \geq 1}$.

Proof. Using the Lipschitzness of ℓ_t and Equation (6.2) we obtain

$$\begin{aligned} & \frac{1}{2}(\ell_t(X_t^+) + \ell_t(X_t^-)) + \tilde{\ell}_t((1 - \xi)x) \\ & \leq \frac{1}{2}(\ell_t(x_t) + \delta G \|S\| + \ell_t(x_t) + \delta G \|S\|) + \ell_t((1 - \xi)x) + \delta G \\ & \leq \ell_t(x_t) + \ell_t(x) + 2\delta G + \xi GR \\ & \leq \tilde{\ell}_t(x_t) + \ell_t(x) + 3\delta G + \xi GR. \end{aligned}$$

In the second step we used $\tilde{\ell}_t((1 - \xi)x) \leq \xi G \|x\| \leq \xi GR$ which results from the Lipschitzness of ℓ_t and the assumption $\mathcal{K} \subseteq R\mathbb{B}$. \square

Next, we show that the second moment of \tilde{g}_t can be controlled by exploiting the Lipschitzness of ℓ_t . In particular,

$$\|\tilde{g}_t(x)\| = \frac{d}{2\delta} |\ell_t(x + \delta S) - \ell_t(x - \delta S)| \|S\| \leq \frac{Gd}{2\delta} \|2\delta S\| = Gd.$$

We are now ready to prove the main result of this section. Namely, that the pseudo-regret of OSGD run using the gradient estimate (Equation (6.1)) is of order \sqrt{n} . We assume that the point \tilde{X}_t played by OSGD at each time t is randomly drawn between the two points X_t^+ and X_t^- , where the loss function is queried.

Theorem 6.4 (Regret of OSGD with two-point feedback). Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a closed convex set such that $r\mathbb{B} \subseteq \mathcal{K} \subseteq R\mathbb{B}$ for some $r, R > 0$. Let \mathcal{L} be a set of G -Lipschitz differentiable and convex losses. Fix $\delta > 0$ and assume OSGD is run on $(1 - \frac{\delta}{r})\mathcal{K}$ with learning rate $\eta > 0$ and gradient estimates (Equation (6.1)),

$$\tilde{g}_t(x_t) = \frac{d}{2\delta} (\ell_t(X_t^+) - \ell_t(X_t^-)) S_t$$

where $S_1, S_2, \dots \in \mathbb{S}$ are independent. For each $t = 1, 2, \dots$ let \tilde{X}_t be drawn at random between X_t^+ and X_t^- . Then the following holds

$$\bar{R}_n \leq \frac{R^2}{\eta} + \eta(Gd)^2 n + \delta \left(3 + \frac{R}{r} \right) Gn.$$

Moreover, if $\eta = \frac{R}{GD\sqrt{n}}$ then for $\delta \rightarrow 0$ we have that

$$\bar{R}_n \leq 2RGd\sqrt{n}.$$

Proof. First of all, we must check that the points $X_t^+ = x_t + \delta S$ and $X_t^- = x_t - \delta S$ on which ℓ_t is queried belong to \mathcal{K} . To see this, recall that $x_t \in (1 - \frac{\delta}{r})\mathcal{K}$. Now, setting $\alpha = \frac{\delta}{r}$ we have that $X_t^+, X_t^- \in (1 - \alpha)\mathcal{K} + \alpha r\mathbb{S}$. Since $r\mathbb{S} \subseteq \mathcal{K}$ and \mathcal{K} is convex, we obtain $(1 - \alpha)\mathcal{K} + \alpha r\mathbb{S} \subseteq (1 - \alpha)\mathcal{K} + \alpha\mathcal{K} \subseteq \mathcal{K}$. Hence, using Lemma 6.3 with the choice $\xi = \frac{\delta}{r}$ we immediately get that for all $x \in \mathcal{K}$,

$$\begin{aligned} & \sum_{t=1}^n \mathbb{E}(\ell_t(\tilde{X}_t) | X_t^+, X_t^-) - \sum_{t=1}^n \ell_t(x) \\ & \leq \frac{1}{2} \sum_{t=1}^n (\ell_t(X_t^+) + \ell_t(X_t^-)) - \sum_{t=1}^n \ell_t(x) \\ & \leq \sum_{t=1}^n \tilde{\ell}_t(x_t) - \sum_{t=1}^n \tilde{\ell}_t\left(\left(1 - \frac{\delta}{r}\right)x\right) + \delta\left(3 + \frac{R}{r}\right)Gn. \end{aligned}$$

Since we already related the loss of \tilde{X}_t to the loss of x_t , we can now apply Theorem 5.7 in the special case of $\tilde{x}_t = x_t$ and with the sequence of losses $(\tilde{\ell}_t)$. This gives

$$\begin{aligned} \mathbb{E} \sum_{t=1}^n \tilde{\ell}_t(x_t) - \mathbb{E} \sum_{t=1}^n \tilde{\ell}_t\left(\left(1 - \frac{\delta}{r}\right)x\right) & \leq \frac{R^2}{\eta} + \eta \sum_{t=1}^n \mathbb{E} \|\tilde{g}_t(x_t)\|^2 \\ & \leq \frac{R^2}{\eta} + \eta(Gd)^2n \end{aligned}$$

where we overapproximated $\|(1 - \frac{\delta}{r})\mathcal{K}\| \leq \|\mathcal{K}\| = R$. This concludes the proof. \square

6.2 One-Point Bandit Feedback

Building on the analysis of the previous section, it is not hard to show that the pseudo-regret can be bounded even when the loss function at each time t is queried in only *one* point. However, we pay this reduced bandit feedback with a worse rate of $n^{3/4}$ in the pseudo-regret bound. It is not known if this rate is optimal, or if it is possible to get a \sqrt{n} regret as in the two-point setting.

The one-point estimate at time t is defined by

$$\tilde{g}_t(x) = \frac{d}{\delta} \ell_t(x + \delta S) S \tag{6.3}$$

where S is drawn at random from \mathbb{S} . Obviously, Lemma 6.2 can be applied to get $\mathbb{E}\tilde{g}_t(x) = \nabla\tilde{\ell}_t(x)$ where, we recall, $\tilde{\ell}_t(x) = \mathbb{E}\ell_t(x + \delta B)$. Differences with the two-point case arise when we bound the second moment of this new \tilde{g}_t . Indeed, if $x + \delta S \in \mathcal{K}$ and the maximum value of each ℓ_t in \mathcal{K} is bounded by L , then

$$\|\tilde{g}_t(x)\| = \frac{d}{\delta}|\ell_t(x + \delta S)|\|S\| \leq \frac{dL}{\delta}.$$

Note the inverse dependence on δ . This dependence plays a key role in the final bound, as the next result shows.

Theorem 6.5 (Regret of OSGD with one-point feedback). Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a closed convex set such that $r\mathbb{B} \subseteq \mathcal{K} \subseteq R\mathbb{B}$ for some $r, R > 0$. Let \mathcal{L} be a set of G -Lipschitz differentiable and convex losses, uniformly bounded by L (that is, $\|\ell\|_\infty \leq L, \forall \ell \in \mathcal{L}$). Fix $\delta > 0$ and assume OSGD is run on $(1 - \frac{\delta}{r})\mathcal{K}$ with learning rate $\eta > 0$ and gradient estimates (Equation (6.3)),

$$\tilde{g}_t(x_t) = \frac{d}{\delta}\ell_t(\tilde{X}_t)S_t$$

where $\tilde{X}_t = x_t + \delta S_t$ and $S_1, S_2, \dots \in \mathbb{S}$ are independent. Then the following holds

$$\bar{R}_n \leq \frac{R^2}{\eta} + \frac{(dL)^2}{\delta^2}\eta n + \delta \left(3 + \frac{R}{r}\right) Gn.$$

Moreover, if

$$\delta = \frac{1}{(2n)^{1/4}} \sqrt{\frac{RdL}{\left(3 + \frac{R}{r}\right)G}} \quad \text{and} \quad \eta = \frac{1}{(2n)^{3/4}} \sqrt{\frac{R^3}{dL\left(3 + \frac{R}{r}\right)G}}$$

then

$$\bar{R}_n \leq 4n^{3/4} \sqrt{RdL \left(3 + \frac{R}{r}\right) G}.$$

Proof. The proof follows along the same lines as the proof of Theorem 6.4. Indeed, we can show that the points $\tilde{X}_t = x_t + \delta S$ on which ℓ_t

is queried belong to \mathcal{K} . Then, using an easy modification of Lemma 6.3 we get that for all $x \in \mathcal{K}$,

$$\begin{aligned} & \sum_{t=1}^n \mathbb{E}(\ell_t(\tilde{X}_t) | X_t^+, X_t^-) - \sum_{t=1}^n \ell_t(x) \\ & \leq \sum_{t=1}^n \tilde{\ell}_t(x_t) - \sum_{t=1}^n \tilde{\ell}_t\left(\left(1 - \frac{\delta}{r}\right)x\right) + \delta\left(3 + \frac{R}{r}\right) Gn. \end{aligned}$$

Applying Theorem 5.7 as in the proof of Theorem 6.4 gives

$$\begin{aligned} \mathbb{E} \sum_{t=1}^n \tilde{\ell}_t(x_t) - \mathbb{E} \sum_{t=1}^n \tilde{\ell}_t\left(\left(1 - \frac{\delta}{r}\right)x\right) & \leq \frac{R^2}{\eta} + \eta \sum_{t=1}^n \mathbb{E} \|\tilde{g}_t(x_t)\|^2 \\ & \leq \frac{R^2}{\eta} + \frac{(dL)^2}{\delta^2} \eta n. \quad \square \end{aligned}$$

6.3 Nonlinear Stochastic Bandits

We conclude with a simple example of nonlinear bandits in the stochastic setting. Unlike the gain-based analysis of stochastic bandits of Section 2, here we keep in with the convention used throughout this section and work exclusively with losses.

We consider a simple unidimensional setting where arms are points in the unit interval $[0, 1]$. If at time t a point $x_t \in [0, 1]$ is played, the loss is the realization of an independent random variable $Y_t \in [0, 1]$ with expected value $\mathbb{E}[Y_t | x_t] = \mu(x_t)$, where $\mu : [0, 1] \rightarrow [0, 1]$ is a fixed but unknown mean loss function. Similarly to Section 2, here the pseudo-regret after n plays of a given strategy can be rewritten as

$$\bar{R}_n = \sum_{t=1}^n \mu(x_t) - n \max_{x \in [0, 1]} \mu(x)$$

where $x_1, \dots, x_n \in [0, 1]$ denote the points played by the strategy.

Throughout this section, we assume that $\mu : [0, 1] \rightarrow [0, 1]$ is unimodal, but not necessarily convex. This means that there exists a unique $x^* = \operatorname{argmin}_{x \in [0, 1]} \mu(x)$ such that $\mu(x)$ is monotone decreasing for $x \in [0, x^*]$ and monotone increasing for $x \in [x^*, 1]$. For example, if μ can be written as $\mu(x) = x f(x)$, where $f : [0, 1] \rightarrow [0, 1]$ is differentiable, monotone decreasing, and such that $x f'(x)$ is strictly decreasing with $f(0) > 0$, then μ is unimodal.

The bandit strategy we analyze in this section is based on the golden section search due to Kiefer [117], which is a general algorithm for finding the extremum of a unimodal function. Similarly to binary search, each step of golden section search narrows the interval in which the extremum is found by querying the function value at certain points that are chosen depending on the outcome of previous queries. Each query shrinks the interval by a factor of $\frac{1}{\varphi} = 0.618\dots$, where $\varphi = \frac{1}{2}(1 + \sqrt{5})$ is the golden ratio.

In our case, queries (i.e., plays) at x return a perturbed version of $\mu(x)$. Since μ is bounded, Hoeffding bounds ensure that we can find the minimum of μ by repeatedly querying each point x requested by the golden search algorithm. However, in order to have a lower bound on the accuracy with which each μ needs to be estimated, we must assume the following condition: there exists $C_L > 0$ such that

$$|\mu(x) - \mu(x')| \geq C_L|x - x'| \quad (6.4)$$

for each x, x' that belong either to $[0, x^* - 1/C_L]$ or to $[x^* + 1/C_L, 1]$.

Finally, irrespective to the uncertainty in the evaluation of μ , in order to bound the regret incurred by golden section search we need a Lipschitz condition on μ . Namely, there exists $C_H > 0$ such that $|\mu(x) - \mu(x')| \leq C_H|x - x'|$ for all $x, x' \in [0, 1]$.

We are now ready to introduce our stochastic version of the golden section search algorithm.

SGS (Stochastic Golden Search):

Parameters: $\varepsilon_1, \varepsilon_2, \dots > 0$.

Initialize: $x_A = 0$ $x_B = \frac{1}{\varphi^2}$ $x_C = 1$.

For each stage $s = 1, \dots, n$

- (1) Let $x'_B = \begin{cases} x_B - \frac{1}{\varphi^2}(x_B - x_A) & x_B - x_A > x_C - x_B \\ x_B + \frac{1}{\varphi^2}(x_C - x_B) & \text{otherwise} \end{cases}$ and
rename points x_B, x'_B so that $x_A < x_B < x'_B < x_C$.
- (2) Play each point in $\{x_A, x_B, x'_B, x_C\}$ for $\frac{2}{\varepsilon_s^2} \ln(6n)$ times and let \hat{x} be the point with lowest total loss in this stage.
- (3) If $\hat{x} \in \{x_A, x_B\}$ then eliminate interval $(x'_B, x_C]$ and let $x_C = x'_B$,
- (4) else eliminate interval $[x_A, x_B)$ and let $x_A = x_B$.

Recall that golden section search proceeds as follows: given three queried points $x_A < x_B < x_C$ where the distance of x_B to the other two points is in the golden ratio (x_B might be closer to x_A or to x_C depending on past queries), the next point x'_B is queried in the largest interval between $x_B - x_A$ and $x_C - x_B$ so that the distance of x'_B to the extrema of that largest interval is in the golden ratio. Assume the resulting ordering is $x_A < x_B < x'_B < x_C$. Then we drop either $[x_A, x_B)$ or $(x'_B, x_C]$ according to whether the smallest value of μ is found in, respectively, $\{x'_B, x_C\}$ or $\{x'_B, x_C\}$. The remaining triplet is such that the distance of the middle point to the other two is again in the golden ratio.

Using elementary algebraic identities for φ , one can show that setting $x_C - x_A = 1$ the following equalities hold at any step of SGS:

$$x_B - x_A = \frac{1}{\varphi^2} \quad x'_B - x_B = \frac{1}{\varphi^3} \quad x_C - x'_B = \frac{1}{\varphi^2}. \quad (6.5)$$

Since either $x_B - x_A$ or $x_C - x'_B$ are eliminated at each stage, at each stage SGS shrinks the search interval by a factor of $1 - \varphi^{-2} = \frac{1}{\varphi}$.

Let $[x_{A,s}, x_{B,s}]$ be the search interval at the beginning of stage $s + 1$, where $x_{A,0} = 0$ and $x_{C,0} = 1$.

Lemma 6.6. If $\varepsilon_s = C_L \varphi^{-(s+3)}$ then

$$\mathbb{P}(x^* \notin [x_{A,s}, x_{C,s}]) \leq \frac{s}{n}$$

holds uniformly over all stages $s \geq 1$.

Proof. Once the interval containing x^* is eliminated it is never recovered, thus we have

$$\begin{aligned} \mathbb{P}(x^* \notin [x_{A,s}, x_{C,s}]) &\leq \mathbb{P}(x^* \notin [x_{A,s-1}, x_{C,s-1}]) \\ &\quad + \mathbb{P}(x^* \notin [x_{A,s}, x_{C,s}] \mid x^* \in [x_{A,s-1}, x_{C,s-1}]). \end{aligned} \quad (6.6)$$

Let $X_s = \{x_{A,s-1}, x_{B,s-1}, x'_{B,s-1}, x_{C,s-1}\}$ where $x_{B,s-1} < x'_{B,s-1}$ are computed in step 1 of stage s . Let $\hat{\mu}_s(x)$ be the sample loss of point $x \in X_s$ in stage s and let $\hat{x}_s = \operatorname{argmin}_{x \in X_s} \hat{\mu}_s(x)$. Since at stage s every point in X_s is played $\frac{2}{\varepsilon_s^2} \ln(6n)$ times,¹ Hoeffding bounds imply that

¹For simplicity, we assume these numbers are integers.

$|\mu(x) - \widehat{\mu}_s(x)| \leq \frac{1}{2}\varepsilon_s$ with probability at least $1 - \frac{1}{6n}$ for all $x \in X_s$ simultaneously. Let

$$x_s^* = \underset{x \in X_s}{\operatorname{argmin}} \mu(x).$$

Now assume $x^* \in [x_{A,s-1}, x_{B,s-1}]$. Then $x^* \notin [x_{A,s}, x_{C,s}]$ implies $\widehat{\mu}_s(x_{B',s-1}) < \widehat{\mu}(x_{B,s-1})$ or $\widehat{\mu}_s(x_{C,s-1}) < \widehat{\mu}(x_{B,s-1})$. Similarly, assume $x^* \in [x_{B',s-1}, x_{C,s-1}]$. Then $x^* \notin [x_{A,s}, x_{C,s}]$ implies $\widehat{\mu}_s(x_{A,s-1}) < \widehat{\mu}(x_{B',s-1})$ or $\widehat{\mu}_s(x_{B,s-1}) < \widehat{\mu}(x_{B',s-1})$. In both cases, we need to compare three values of μ on the same side with respect to x^* . (When $x^* \in [x_{B,s-1}, x_{B',s-1}]$ we always have $x^* \in [x_{A,s}, x_{C,s}]$.) Hence, we can apply our assumption involving C_L . More precisely, Equation (6.5) implies that after s stages the search interval has size φ^{-s} and $\min\{x_{B,s} - x_{A,s}, x'_{B,s} - x_{B,s}, x_{C,s} - x'_{B,s}\} = \varphi^{-(s+3)}$. Hence, introducing

$$\Delta_s = \min\{|\mu(x_{B,s}) - \mu(x_{A,s})|, |\mu(x'_{B,s}) - \mu(x_{B,s})|, |\mu(x_{C,s}) - \mu(x'_{B,s})|\}$$

we have

$$\Delta_s \geq C_L \min\{x_{B,s} - x_{A,s}, x'_{B,s} - x_{B,s}, x_{C,s} - x'_{B,s}\} \geq C_L \varphi^{-(s+3)} = \varepsilon_s.$$

Let $T_s = \frac{8}{\varepsilon_s^2} \ln(6n)$ the length of stage s . We can write

$$\begin{aligned} & \mathbb{P}(x^* \notin [x_{A,s}, x_{C,s}] | x^* \in [x_{A,s-1}, x_{C,s-1}]) \\ &= \mathbb{P}(\widehat{\mu}_s(\widehat{x}_s) < \widehat{\mu}(x_s^*)) \\ &\leq \sum_{x \in X_s \setminus \{x_s^*\}} \mathbb{P}(\widehat{\mu}_s(x) < \widehat{\mu}(x_s^*)) \\ &\leq \sum_{x \in X_s \setminus \{x_s^*\}} \left(\mathbb{P}\left(\widehat{\mu}_s(x) < \mu(x) - \frac{\Delta_s}{2}\right) \right. \\ &\quad \left. + \mathbb{P}\left(\mu(x_s^*) < \widehat{\mu}(x_s^*) - \frac{\Delta_s}{2}\right) \right) \\ &\leq 6e^{-T_s \Delta_s^2 / 8} \\ &\leq 6e^{-T_s \varepsilon_s^2 / 8} \leq \frac{1}{n}. \end{aligned}$$

Substituting this in Equation (6.6) and recurring gives the desired result. \square

Theorem 6.7 (Regret of SGS). For any unimodal and C_H -Lipschitz mean loss function $\mu : [0, 1] \rightarrow [0, 1]$ that satisfies Equation (6.4), if the SGS algorithm is run with $\varepsilon_s = C_L \varphi^{-(s+3)}$ then

$$\bar{R}_n \leq \frac{C_H}{C_L^2} 8\varphi^6 \ln(6n) \left[\frac{2\varphi}{\varphi - 1} \sqrt{1 + C_L^2 n} + \frac{1}{4} \log_\varphi^2(1 + C_L^2 n) \right].$$

Proof. We start by decomposing the pseudo-regret as follows,

$$\bar{R}_n \leq \sum_{s=1}^S T_s (\min_{x \in A_s} \mu(x) - \mu(x^*)) + \sum_{s=1}^S \left(\sum_{t \in T_s} \mu(x_t) - T_s \min_{x \in A_s} \mu(x) \right).$$

Using the Lipschitz assumption

$$\max_{x, x' \in A_s} |\mu(x) - \mu(x')| \leq C_H |x_{C,s} - x_{A,s}|$$

and recalling that $|x_{C,s} - x_{A,s}| \leq \varphi^{-s}$, we bound the first term in this decomposition as follows

$$\begin{aligned} \sum_{s=1}^S T_s \left(\min_{x \in A_s} \mu(x) - \mu(x^*) \right) &\leq T_s C_H |x_{C,s} - x_{A,s}| \mathbb{P}(x^* \in [x_{A,s}, x_{C,s}]) \\ &\quad + T_s C_H \mathbb{P}(x^* \notin [x_{A,s}, x_{C,s}]) \\ &\leq \frac{T_s C_H}{\varphi^s} + T_s C_H \frac{s}{n}. \end{aligned}$$

The second term is controlled similarly,

$$\sum_{s=1}^S \left(\sum_{t \in T_s} \mu(x_t) - T_s \min_{x \in A_s} \mu(x) \right) \leq T_s C_H |x_{C,s} - x_{A,s}| \leq \frac{T_s C_H}{\varphi^s}.$$

Hence we get an easy expression for the regret,

$$\begin{aligned} \bar{R}_n &\leq C_H \sum_{s=1}^S T_s \left(\frac{2}{\varphi^s} + \frac{s}{n} \right) \\ &\leq \frac{C_H}{C_L^2} 8\varphi^6 \ln(6n) \sum_{s=1}^S \varphi^{2s} \left(\frac{2}{\varphi^s} + \frac{s}{n} \right). \end{aligned} \quad (6.7)$$

We now compute an upper bound on the number S of stages,

$$n \leq \sum_{s=1}^S T_s = \frac{8\varphi^6}{C_L^2} \ln(6n) \sum_{s=1}^S \varphi^{2s} \leq \frac{8\varphi^6}{C_L^2} \ln(6n) \frac{\varphi^{2S+2}}{\varphi^2 - 1}.$$

Solving for n and overapproximating we get

$$S \leq \frac{1}{2} \log_{\varphi}(1 + C_L^2 n).$$

Therefore, the sum in Equation (6.7) is bounded as follows

$$\begin{aligned} 2 \sum_{s=1}^S \varphi^s + S^2 &\leq \frac{2\varphi}{\varphi - 1} \varphi^S + S^2 \\ &\leq \frac{2\varphi}{\varphi - 1} \sqrt{1 + C_L^2 n} + \frac{1}{4} \log_{\varphi}^2(1 + C_L^2 n). \end{aligned}$$

Substituting the above in Equation (6.7) concludes the proof. \square

6.4 Bibliographic Remarks

Gradient-free methods for stochastic approximation have been studied for a long time — see the bibliographic remarks at the end of Section 5 for some references. However, relatively few results provide regret bounds. The approach presented in this section for online convex optimization with bandit feedback was pioneered by Flaxman et al. [79] and Kleinberg [120]. The improved rate for the two-point model was later shown in [6].

While the results for nonlinear bandits in the adversarial model are still scarce, there is a far richer body of work in the stochastic model. The result based on golden section search presented in Section 6.3 is due to Yu and Mannor [171]. It represents only a tiny portion of the known results in the stochastic model. In the general case of Lipschitz mean-payoff on a compact subset of \mathbb{R}^d , it can be shown that the minimax regret is $\Omega(n^{\frac{d+1}{d+2}})$. Thus, the rate rapidly deteriorates as the dimension increases, a phenomenon known as the *curse of dimensionality*. However it was shown in [51, 122] that under a generalized version of Equation (6.4) it is possible to circumvent the curse of dimensionality and obtain a regret of $\tilde{O}(\sqrt{n})$. This result builds upon and generalizes

a sequence of works that include the discretization approach (for the one-dimensional case) of Kleinberg [120] and Auer et al. [26], as well as the method of Cope [68] based on the Kiefer–Wolfowitz procedure (a classical method of stochastic optimization). The key new algorithmic idea introduced in [51, 122] is to adaptively partition the set of actions in order to exploit the smoothness of the mean-payoff function around its maximum. We refer the reader to [52] for the details of this result (which is much more general than what we briefly outlined, in particular it applies for metric spaces, or even more general action sets), as well as a more precise historical account.

Another direction for nonlinear stochastic bandits was recently investigated in [8]. In this work the authors study the case of a convex mean loss function, and they show how to combine the zeroth-order optimization method of Nemirovski and Yudin [139] with a “center point device” to obtain a regret of $\tilde{O}(\sqrt{n})$.

A more general version of nonlinear stochastic bandit was also studied in [12]. In this paper the authors assume that the mean-payoff function lies in some known set of functions \mathcal{F} . They define a notion of complexity for the class \mathcal{F} , the haystack dimension $\text{HD}(\mathcal{F})$, and they show that the worst case regret in \mathcal{F} is lower bounded by $\text{HD}(\mathcal{F})$. Unfortunately their upper bound does not match the lower bound, and the authors suggest that the definition of the haystack dimension should be modified in order to obtain matching upper and lower bounds.

Finally, a related problem in a Bayesian setting was studied in [89, 155], where it is assumed that the payoff functions are drawn from Gaussian processes.

7

Variants

In the previous sections we explored a few fundamental variations around the basic multi-armed bandit problem. In both the stochastic and adversarial frameworks, these variants basically revolved around a single principle: by adding constraints on the losses (or rewards), it is possible to compete against larger sets of arms. While this is indeed a fundamental axis in the space of bandit problems, it is important to realize that there are many other directions. Indeed, we might sketch a “bandit space” spanning the following coordinates:

- Evolution of payoffs over time: stochastic, adversarial, Markovian, etc.
- Structure of payoff functions: linear, Lipschitz, Gaussian process, etc.
- Feedback structure: full information, bandit, semi-bandit, partial monitoring, etc.
- Context structure (if any).
- Notion of regret.

Clearly, such extensions greatly increase the number of potential applications of bandit models. While many of these extensions were already

discussed in the previous sections, in the following we focus on others (such as the sleeping bandits or the truthful bandits) so to visit more exotic regions of the bandit space.

7.1 Markov Decision Processes, Restless and Sleeping Bandits

Extending further the model of Markovian bandits (mentioned at the end of Section 1), one can also define a general Markov Decision Process (MDP) — see also Section 7.1. For example, the stochastic bandit of Section 2 corresponds to a single-state MDP.

In full generality, a finite MDP can be described by a set of states $\{1, \dots, S\}$, a set of actions $\{1, \dots, K\}$, a set $\{p_{i,s}, 1 \leq i \leq K, 1 \leq s \leq S\}$ of transition distributions over S , and a set $\{\nu_{i,s}, 1 \leq i \leq K, 1 \leq s \leq S\}$ of reward distributions over $[0, 1]$. In this model, taking action i in state s generates a stochastic reward drawn from $\nu_{i,s}$ and a Markovian transition to a state drawn from $p_{i,s}$. Similarly to the multi-armed bandit problem, here one typically assumes that the reward distributions and transition distributions are unknown, and the goal is to navigate through the MDP so as to maximize some function of the obtained rewards. The field that studies this type of problem is called Reinforcement Learning. The interested reader is addressed to [108, 157, 158]. Reinforcement learning results with a flavor similar to those described in the previous sections can be found in [48, 105, 141, 172].

An intermediate model, between stochastic multi-armed bandits and MDPs, is the one of restless bandits. As in Markovian bandits, each arm is associated with a Markovian reward process with its own state space. Each time an arm is chosen, the associated Markov process generates an observable reward and makes a transition to a new state, which is also observed. However, unlike Markovian bandits an unobserved transition occurs for each arm that is not chosen. Using concentration inequalities for Markov chains — see, e.g., [127], one can basically show that, under suitable assumptions, UCB attains a logarithmic regret for restless bandits as well — see [77, 160]. A more general regret bound for restless bandits has been recently proven by Ortner et al. [142].

An apparently similar problem was studied by Garivier and Moulines [84], where they assume that the reward distributions can abruptly change at unknown time instants (and there is a small number of such change-points). Within this model, the authors prove that the best possible regret is of order \sqrt{n} , which is matched by the Exp3.P algorithm — see the discussion in Section 3.4.3. Thus, while the two problems (restless bandits and bandits with change-points) might look similar, they are fundamentally different. In particular, note that the latter problem cannot be cast as an MDP.

Another intermediate model, with important applications, is that of the sleeping bandits. There, it is assumed that the set of available actions is varying over time. We refer the interested reader to [113, 112, 121, 152] for the details of this model as well as the theoretical guarantees that can be obtained. A somewhat related problem was also studied in [90] where it is assumed that the set of arms becomes unavailable for a random time after each arm pull (and the distribution of this random time depends on the selected arm).

7.2 Pure Exploration Problems

The focus of bandits, and most of their variants, is on problems where there is a notion of cumulative rewards, which is to be maximized. This criterion leaves out a number of important applications where there is an online aspect (e.g., sequential decisions), but the goal is not maximizing cumulative rewards. The simplest example is perhaps the pure exploration version of stochastic bandits. In this model, at the end of round n the algorithm has to output a recommendation J_n which represents its estimate for the optimal arm. The focus here is on the control of the so-called simple regret, introduced by Bubeck et al. [49, 50] and defined as $r_n = \mu^* - \mathbb{E}\mu_{J_n}$.

Bubeck et al. [49] prove that minimizing the simple regret is fundamentally different from minimizing the pseudo-regret \bar{R}_n , in the sense that one always have $r_n \geq \exp(-C\bar{R}_n)$ for some constant $C > 0$ (which depends on the reward distributions). Thus, this regret calls for different bandit algorithms. Audibert et al. [20] exhibit a simple strategy with optimal performances up to a logarithmic factor. The idea

is very simple: the strategy SR (Successive Rejects) works in $K - 1$ phases. SR keeps a set of active arms that are sampled uniformly in each phase. At the end of a phase, the arm with smallest empirical mean is removed from the set of active arms. It can be shown that this strategy has a simple regret of order $\exp(-c \frac{n}{H \ln K})$, where $H = \sum_{i \neq i^*} \frac{1}{\Delta_i^2}$ is the complexity measure of identifying the best arm, and c is a numerical constant. Moreover, a matching lower bound (up to logarithmic factors) was also proved. These ideas were extended in different ways by Gabillon et al. [82], Bui et al. [55] and Bubeck et al. [54].

A similar problem was studied in a PAC model by Even-Dar et al. [75]. The goal is to find, with probability at least $1 - \delta$, an arm with mean at least ε close the optimal mean, and the relevant quantity is the number of pulls needed to achieve this goal. For this problem, the authors derive an algorithm called Successive Elimination that achieves an optimal number of pulls (up to logarithmic factors). Successive Elimination works as follows: it keeps an estimate of the mean of each arm, together with a confidence interval. If two confidence intervals are disjoint, then the arm with the lowest confidence interval is eliminated. Using this procedure, one can achieve the (ε, δ) guarantee with a number of pulls of order $H \ln \frac{K}{\Delta}$. A matching lower bound is due to Mannor and Tsitsiklis [134], and further results are discussed by Even-Dar et al. [76].

In some applications one is not interested in the best arm, but rather in having a good estimate of the mean μ_i for each arm. In this setting a reasonable measure of performance is given by

$$L_n = \mathbb{E} \sum_{i=1}^K (\mu_i - \hat{\mu}_{i, T_i(n)})^2.$$

Clearly, the optimal static allocation depends only on the variances of the arms, and we denote by L_n^* the performance of this strategy. This setting was introduced by Antos et al. [13], where the authors studied the regret $L_n - L_n^*$, and showed that a regret of order $n^{-3/2}$ was achievable. This result was then refined by Carpentier et al. [59] and Carpentier and Munos [60]. The basic idea in these papers is to resort to the optimism in face of uncertainty principle, and to approximate

the optimal static allocation by replacing the true variance with an upper confidence bound on it.

7.3 Dueling Bandits

An interesting variation of stochastic bandits was recently studied by Yue et al. [173]. The model considered in this paper is called dueling bandits. The main idea is that the player has to choose a pair of arms (I_t, J_t) at each round and can only observe the relative performances of these two arms, i.e., the player only knows which arm had the highest reward. More formally, in dueling bandits we assume that there exists a total ordering \succ on $\{1, \dots, K\}$ with the following properties:

- (1) If $i \succ j$, then the probability that the reward of arm i is larger than the reward of arm j is equal to $\frac{1}{2} + \Delta_{i,j}$ with $\Delta_{i,j} > 0$.
- (2) If $i \succ j \succ k$, then $\Delta_{i,j} + \Delta_{j,k} \geq \Delta_{i,k} \geq \max\{\Delta_{i,j}, \Delta_{j,k}\}$.

Upon selecting a pair (I_t, J_t) , the player receives a random variable drawn from a Bernoulli distribution with parameter $\frac{1}{2} + \Delta_{i,j}$. In this setting a natural regret notion is the following quantity, where i^* is the largest element in the ordering \succ ,

$$\mathbb{E} \sum_{t=1}^n (\Delta_{i^*, I_t} + \Delta_{i^*, J_t}).$$

It was proved in Yue et al. [173] that the optimal regret for this problem is of order $\frac{K}{\Delta} \log n$, where $\Delta = \min_{i \neq j} \Delta_{i,j}$. A simple strategy that attains this rate, based on the Successive Elimination algorithm of Even-Dar et al. [75], was proposed by Yue and Joachims [174].

7.4 Discovery with Probabilistic Expert Advice

Bubeck et al. [47] study a model with a stochastic bandit flavor (in fact it can be cast as an MDP), where the key for the analysis is a sort of “nonlinear” regret bound. In this model rewards represent items in some set \mathcal{X} which is partitioned in a subset $A \subset \mathcal{X}$ of interesting items and in a subset $\mathcal{X} \setminus A$ of noninteresting items. The goal is to maximize the total expected number of interesting items found after n pulls,

where observing twice the same item does not help. A natural notion of regret is obtained by comparing the number of interesting items $F(n)$ found by a given strategy to the number $F^*(n)$ found by the optimal strategy. It turns out that analyzing such regret directly is difficult. The first issue is that in this problem the notion of a “good” arm is dynamic, in the sense that an arm could be very good for a few pulls and then completely useless. Furthermore, a strategy making bad decisions in the beginning will have better opportunities in the future than the optimal strategy (which already discovered some interesting items). Taking into account these issues, it turns out that it is easier to show that for good strategies, $F(n)$ is not too far from $F^*(n')$, where n' is not much smaller than n . Such a statement — which can be interpreted as a nonlinear regret bound — shows that the analyzed strategy slightly “lags” behind the optimal strategy. In Bubeck et al. [47] a nonlinear regret bound is derived for an algorithm based on estimating the mass of interesting items left on each arm (the so-called Good-Turing estimator), combined with the optimism in face of uncertainty principle of Section 2. We refer the reader to Bubeck et al. [47] for more precise statements.

7.5 Many-Armed Bandits

The many-armed bandit setting was introduced by Berry et al. [38], and then extended and refined by Wang et al. [167]. This setting corresponds to a stochastic bandit with an infinite number of arms. The extra assumption that makes this problem feasible is a prior knowledge on the distribution of the arms. More precisely, when the player asks to “add” a new arm to his current set of active arms, one assumes that the probability that this arm is ε -optimal is of order ε^β , for some known $\beta > 0$. Thus, the player faces a trade-off between exploitation, exploration, and discovery, where the last component comes from the fact that the player needs to consider new arms to make sure that he has an active ε -optimal arm. Using a UCB strategy on the active arms, and by adding new arms at a rate which depends on β , Wang et al. [167] prove that a regret of order

$$n^{\max\{\frac{1}{2}, \frac{\beta}{1+\beta}\}}$$

is achievable in this setting.

7.6 Truthful Bandits

A popular application domain for bandit algorithms is ad placement on the Web. In the pay-per-click model, for each incoming user $t = 1, 2, \dots$ the publisher selects an advertiser I_t from a pool of K advertisers, and display the corresponding ad to the user. The publisher then gets a reward if the ad is clicked by the user. This problem is well modeled by the multi-armed bandit setting. However, there is a fundamental aspect of the ad placement process which is overlooked by this formulation. Indeed, prior to running an ad-selection algorithm (i.e., a bandit algorithm), each advertiser $i \in \{1, \dots, K\}$ issues a bet b_i . This number is how much i is willing to pay for a click. Each bidder keeps also a private value v_i , which is the true value i assigns to a click. Because a rational bidder ensures that $b_i \leq v_i$, the difference $v_i - b_i$ defines the utility for bidder i . The basic idea of *truthful* bandits is to construct a bandit algorithm such that each advertiser has no incentive in submitting a bet b_i such that $b_i < v_i$. A natural question to ask is whether this restriction to truthful algorithms changes the dynamics of the multi-armed bandit problem. This has been investigated in a number of papers, including [29, 30, 73, 170]. Truthful bandits are part of a more general thread of research at the interface between bandits and Mechanism Design.

7.7 Concluding Remarks

As pointed out in the Introduction, the growing interest for bandits arises from the large number of industrially relevant problems that can be modeled as a multi-armed bandit. In particular, the sequential nature of the bandit setting makes it perfectly suited to various Internet and Web applications. These include search engine optimization with dueling bandits, or ad placement with contextual bandits and truthful bandits, see the references in, respectively, Sections 7.3, 4.5, and 7.6.

Multi-armed bandits also proved to be very useful in other areas. For example, thanks to the strong connections between bandits and Markov Decision Processes, a breakthrough in Monte Carlo Tree Search (MCTS) was achieved using bandits ideas. More precisely, based on the

sparse planning idea of Kearns et al. [116], Kocsis and Szepesvári [123] introduced a new MCTS strategy called UCT (UCB applied to Trees) that led to a substantial advancement in Computer Go performance, see Gelly et al. [85]. Note that, from a theoretical point of view UCT was proved to perform poorly by Coquelin and Munos [69], and a strategy based on a similar idea, but with improved theoretical performance, was proposed by Bubeck and Munos [48]. Other applications in related directions have also been explored, see for example [104, 161] and many others.

Many new domains of application for bandits problems are currently investigated. For example: multichannel opportunistic communications, Liu et al. [129]; model selection, Agarwal et al. [7]; boosting, Busa-Fekete and Kegl [57], management of dark pools of liquidity (a recent type of stock exchange), Agarwal et al. [5]; and security analysis of power systems, Bubeck et al. [47].

Given the fast pace of new variants, extensions, and applications coming out every week, we had to make tough decisions about what to present in this monograph. We apologize for everything we had to leave out. On the other hand, we do hope that what we decided to put in will enthruse more researchers about entering this exciting field.

Acknowledgments

We would like to thank Mike Jordan for proposing to write this monograph and James Finlay for keeping us on track. The table of contents was laid down with the help of Gábor Lugosi. We would also like to express our warm thanks to the reviewers, whose insightful comments have led to many substantial improvements. And a special thank goes to the “nonanonymous” referee Alexandrs Slivkins, whose review stands out for breadth and depth. Thank you Alexandrs, you have done a really exceptional job!

Nicolò Cesa-Bianchi gratefully acknowledges partial support by the PASCAL2 Network of Excellence under EC grant no. 216886.

References

- [1] Y. Abbasi-Yadkori, D. Pal, and C. Szepesvári, “Improved algorithms for linear stochastic bandits,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [2] N. Abe and P. M. Long, “Associative reinforcement learning using linear probabilistic concepts,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 1999.
- [3] J. Abernethy, E. Hazan, and A. Rakhlin, “Competing in the dark: An efficient algorithm for bandit linear optimization,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2008.
- [4] J. Abernethy and A. Rakhlin, “Beating the adaptive bandit with high probability,” in *In Proceedings of the Annual Conference on Learning Theory (COLT)*, 2009.
- [5] A. Agarwal, P. Bartlett, and M. Dama, “Optimal allocation strategies for the dark pool problem,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 9, 2010.
- [6] A. Agarwal, O. Dekel, and L. Xiao, “Optimal algorithms for online convex optimization with multi-point bandit feedback,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2010.
- [7] A. Agarwal, J. Duchi, P. L. Bartlett, and C. Levrard, “Oracle inequalities for computationally budgeted model selection,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 19, 2011.

- [8] A. Agarwal, D. P. Foster, D. Hsu, S. M. Kakade, and A. Rakhlin, “Stochastic convex optimization with bandit feedback,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [9] R. Agrawal, “Sample mean based index policies with $\mathcal{O}(\log n)$ regret for the multi-armed bandit problem,” *Advances in Applied Mathematics*, vol. 27, pp. 1054–1078, 1995.
- [10] S. Agrawal and N. Goyal, “Analysis of Thompson sampling for the multi-armed bandit problem,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 23, 2012.
- [11] C. Allenberg, P. Auer, L. Györfi, and G. Ottucsák, “Hannan consistency in on-line learning in case of unbounded losses under partial monitoring,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2006.
- [12] K. Amin, M. Kearns, and U. Syed, “Bandits, query learning, and the haystack dimension,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 19, 2011.
- [13] A. Antos, V. Grover, and C. Szepesvári, “Active learning in multi-armed bandits,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2008.
- [14] R. Arora, O. Dekel, and A. Tewari, “Online bandit learning against an adaptive adversary: From regret to policy regret,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2012.
- [15] S. Arora, E. Hazan, and S. Kale, “The multiplicative weights update method: A meta-algorithm and applications,” *Theory of Computing*, vol. 8, pp. 121–164, 2012.
- [16] K. J. Arrow, D. Blackwell, and M. A. Girshick, “Bayes and minimax solutions of sequential decision problems,” *Econometrica*, pp. 213–244, 1949.
- [17] J.-Y. Audibert and S. Bubeck, “Minimax policies for adversarial and stochastic bandits,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2009.
- [18] J.-Y. Audibert and S. Bubeck, “Regret bounds and minimax policies under partial monitoring,” *Journal of Machine Learning Research*, vol. 11, pp. 2635–2686, 2010.
- [19] J.-Y. Audibert, S. Bubeck, and G. Lugosi, “Minimax policies for combinatorial prediction games,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 19, 2011.
- [20] J.-Y. Audibert, S. Bubeck, and R. Munos, “Best arm identification in multi-armed bandits,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2010.
- [21] J.-Y. Audibert, R. Munos, and C. Szepesvári, “Exploration-exploitation trade-off using variance estimates in multi-armed bandits,” *Theoretical Computer Science*, vol. 410, pp. 1876–1902, 2009.
- [22] P. Auer, “Using confidence bounds for exploitation-exploration trade-offs,” *Journal of Machine Learning Research*, vol. 3, pp. 397–422, 2002.

- [23] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multi-armed bandit problem,” *Machine Learning Journal*, vol. 47, no. 2–3, pp. 235–256, 2002.
- [24] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire, “The non-stochastic multi-armed bandit problem,” *SIAM Journal on Computing*, vol. 32, no. 1, pp. 48–77, 2002.
- [25] P. Auer and R. Ortner, “UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem,” *Periodica Mathematica Hungarica*, vol. 61, pp. 55–65, 2010.
- [26] P. Auer, R. Ortner, and C. Szepesvári, “Improved rates for the stochastic continuum-armed bandit problem,” in *Proceedings of the International Conference on Learning Theory (COLT)*, 2007.
- [27] B. Awerbuch and R. Kleinberg, “Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches,” in *Proceedings of the Annual ACM Symposium on Theory of Computing (STOC)*, 2004.
- [28] A. Baños, “On pseudo-games,” *Annals of Mathematical Statistics*, vol. 39, pp. 1932–1945, 1968.
- [29] M. Babaioff, R. D. Kleinberg, and A. Slivkins, “Truthful mechanisms with implicit payment computation,” in *ACM Conference on Electronic Commerce 2010 (EC)*, 2010.
- [30] M. Babaioff, Y. Sharma, and A. Slivkins, “Characterizing truthful multi-armed bandit mechanisms,” in *ACM Conference on Electronic Commerce 2009 (EC)*, 2009.
- [31] F. Bach and E. Moulines, “Non-asymptotic analysis of stochastic approximation algorithms for machine learning,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [32] K. Ball, “An elementary introduction to modern convex geometry,” in *Flavors of Geometry*, (S. Levy, ed.), pp. 1–58, Cambridge University Press, 1997.
- [33] P. Bartlett, V. Dani, T. Hayes, S. Kakade, A. Rakhlin, and A. Tewari, “High probability regret bounds for online optimization,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2008.
- [34] G. Bartok, D. Pal, and C. Szepesvári, “Toward a classification of finite partial-monitoring games,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2010.
- [35] G. Bartok, D. Pal, C. Szepesvári, and I. Szita, “Online learning,” *Lecture Notes*, 2011.
- [36] A. Beck and M. Teboulle, “Mirror Descent and nonlinear projected subgradient methods for convex optimization,” *Operations Research Letters*, vol. 31, no. 3, pp. 167–175, 2003.
- [37] A. Ben-Tal and A. Nemirovski, “The conjugate barrier Mirror Descent method for non-smooth convex optimization,” Technical Report, MINERVA Optimization Center Report, Faculty of Industrial Engineering and Management, Technion–Israel Institute of Technology, Haifa, 1999.
- [38] D. A. Berry, R. W. Chen, D. C. Heath, A. Zame, and L. A. Shepp, “Bandit problems with infinitely many arms,” *Annals of Statistics*, vol. 25, pp. 2103–2116, 1997.

- [39] A. Beygelzimer, J. Langford, L. Li, L. Reyzin, and R. Schapire, “Contextual bandit algorithms with supervised learning guarantees,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 15, 2011.
- [40] A. Beygelzimer, J. Langford, L. Li, L. Reyzin, and R. E. Schapire, “Contextual bandit algorithms with supervised learning guarantees,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 15, 2011.
- [41] S. Boucheron, O. Bousquet, and G. Lugosi, “Theory of classification: A survey of recent advances,” *ESAIM: Probability and Statistics*, vol. 9, pp. 323–375, 2005.
- [42] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [43] S. Bubeck, “Bandits games and clustering foundations,” PhD thesis, Université Lille 1, 2010.
- [44] S. Bubeck, “Introduction to online optimization,” *Lecture Notes*, 2011.
- [45] S. Bubeck, N. Cesa-Bianchi, and S. M. Kakade, “Towards minimax policies for online linear optimization with bandit feedback,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 23, 2012.
- [46] S. Bubeck, N. Cesa-Bianchi, and G. Lugosi, “Bandits with heavy tail,” Arxiv preprint arXiv:1209.1727, 2012.
- [47] S. Bubeck, D. Ernst, and A. Garivier, “Optimal discovery with probabilistic expert advice,” Arxiv preprint arXiv:1110.5447, 2011.
- [48] S. Bubeck and R. Munos, “Open loop optimistic planning,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2010.
- [49] S. Bubeck, R. Munos, and G. Stoltz, “Pure exploration in multi-armed bandits problems,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2009.
- [50] S. Bubeck, R. Munos, and G. Stoltz, “Pure exploration in finitely-armed and continuously-armed bandits,” *Theoretical Computer Science*, vol. 412, pp. 1832–1852, 2011.
- [51] S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári, “Online optimization in \mathcal{X} -armed bandits,” in *Advances in Neural Information Processing Systems (NIPS)*, 2009.
- [52] S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári, “ \mathcal{X} -armed bandits,” *Journal of Machine Learning Research*, vol. 12, pp. 1587–1627, 2011.
- [53] S. Bubeck and A. Slivkins, “The best of both worlds: Stochastic and adversarial bandits,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 23, 2012.
- [54] S. Bubeck, T. Wang, and N. Viswanathan, “Multiple identifications in multi-armed bandits,” Arxiv preprint arXiv:1205.3181, 2012.
- [55] L. Bui, R. Johari, and S. Mannor, “Committing bandits,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [56] A. N. Burnetas and M. N. Katehakis, “Optimal adaptive policies for Markov decision processes,” *Mathematics of Operations Research*, pp. 222–255, 1997.

- [57] R. Busa-Fekete and B. Kegl, “Fast boosting using adversarial bandits,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2011.
- [58] O. Cappé, A. Garivier, O. Maillard, R. Munos, and G. Stoltz, “Kullback-Leibler upper confidence bounds for optimal sequential allocation,” Arxiv preprint arXiv:1210.1136, 2012.
- [59] A. Carpentier, A. Lazaric, M. Ghavamzadeh, R. Munos, and P. Auer, “Upper confidence bounds algorithms for active learning in multi-armed bandits,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2011.
- [60] A. Carpentier and R. Munos, “Finite time analysis of stratified sampling for Monte Carlo,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [61] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [62] N. Cesa-Bianchi and G. Lugosi, “Combinatorial bandits,” *Journal of Computer and System Sciences*, vol. 78, no. 5, pp. 1404–1422, 2012.
- [63] N. Cesa-Bianchi, G. Lugosi, and G. Stoltz, “Minimizing regret with label efficient prediction,” *IEEE Transactions on Information Theory*, vol. 51, pp. 2152–2162, 2005.
- [64] N. Cesa-Bianchi, Y. Mansour, and G. Stoltz, “Improved second-order bounds for prediction with expert advice,” *Machine Learning*, vol. 66, pp. 321–352, 2007.
- [65] O. Chapelle and L. Li, “An empirical evaluation of Thompson sampling,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [66] W. Chu, L. Li, L. Reyzin, and R. Schapire, “Contextual bandits with linear payoff functions,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 15, 2011.
- [67] A. Conn, K. Scheinberg, and L. Vicente, “Introduction to derivative-free optimization,” in *Society for Industrial and Applied Mathematics (SIAM)*, 2009.
- [68] E. W. Cope, “Regret and convergence bounds for a class of continuum-armed bandit problems,” *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1243–1253, 2009.
- [69] P.-A. Coquelin and R. Munos, “Bandit algorithms for tree search,” in *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2007.
- [70] K. Crammer and C. Gentile, “Multiclass classification with bandit feedback using adaptive regularization,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2011.
- [71] V. Dani, T. Hayes, and S. Kakade, “The price of bandit information for online optimization,” in *Advances in Neural Information Processing Systems (NIPS)*, 2008.
- [72] V. Dani, T. Hayes, and S. Kakade, “Stochastic linear optimization under bandit feedback,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2008.

- [73] N. Devanur and S. M. Kakade, “The price of truthfulness for pay-per-click auctions,” in *ACM Conference on Electronic Commerce 2009 (EC)*, 2009.
- [74] M. Dudik, D. Hsu, S. Kale, N. Karampatziakis, J. Langford, L. Reyzin, and T. Zhang, “Efficient optimal learning for contextual bandits,” in *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2011.
- [75] E. Even-Dar, S. Mannor, and Y. Mansour, “PAC bounds for multi-armed bandit and Markov decision processes,” in *Proceedings of the Annual Conference on Computational Learning Theory (COLT)*, 2002.
- [76] E. Even-Dar, S. Mannor, and Y. Mansour, “Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems,” *Journal of Machine Learning Research*, vol. 7, pp. 1079–1105, 2006.
- [77] S. Filippi, O. Cappé, and A. Garivier, “Optimally sensing a single channel without prior information: The tiling algorithm and regret bounds,” *Selected Topics in Signal Processing*, vol. 5, no. 1, pp. 68–76, 2011.
- [78] S. Filippi, O. Cappé, A. Garivier, and C. Szepesvári, “Parametric bandits: The generalized linear case,” in *Neural Information Processing Systems (NIPS)*, 2010.
- [79] A. Flaxman, A. Kalai, and B. McMahan, “Online convex optimization in the bandit setting: Gradient descent without a gradient,” in *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2005.
- [80] D. Foster and A. Rakhlin, “No internal regret via neighborhood watch,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 22, 2012.
- [81] D. Foster and R. Vohra, “Asymptotic calibration,” *Biometrika*, vol. 85, pp. 379–390, 1998.
- [82] V. Gabillon, M. Ghavamzadeh, A. Lazaric, and S. Bubeck, “Multi-bandit best arm identification,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [83] A. Garivier and O. Cappé, “The KL-UCB algorithm for bounded stochastic bandits and beyond,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 19, 2011.
- [84] A. Garivier and E. Moulines, “On upper-confidence bound policies for switching bandit problems,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2011.
- [85] S. Gelly, Y. Wang, R. Munos, and O. Teytaud, “Modification of UCT with patterns in Monte-Carlo Go,” Technical Report RR-6062, INRIA, 2006.
- [86] J. Gittins, K. Glazebrook, and R. Weber, *Multi-armed Bandit Allocation Indices*. John Wiley and Sons, 2nd ed., 2011.
- [87] J. C. Gittins, “Bandit processes and dynamic allocation indices,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 148–177, 1979.
- [88] A. Grove, N. Littlestone, and D. Schuurmans, “General convergence results for linear discriminant updates,” *Machine Learning*, vol. 43, pp. 173–210, 2001.

- [89] S. Grünewälder, J.-Y. Audibert, M. Opper, and J. Shawe-Taylor, “Regret bounds for Gaussian process bandit problems,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 9, 2010.
- [90] A. Györfy, L. Kocsis, I. Szabó, and C. Szepesvári, “Continuous time associative bandit problems,” in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
- [91] A. Györfy, T. Linder, G. Lugosi, and G. Ottucsák, “The on-line shortest path problem under partial monitoring,” *Journal of Machine Learning Research*, vol. 8, pp. 2369–2403, 2007.
- [92] J. Hannan, “Approximation to Bayes risk in repeated play,” *Contributions to the Theory of Games*, vol. 3, pp. 97–139, 1957.
- [93] S. Hart and A. Mas-Colell, “A simple adaptive procedure leading to correlated equilibrium,” *Econometrica*, vol. 68, pp. 1127–1150, 2000.
- [94] S. Hart and A. Mas-Colell, “A general class of adaptive strategies,” *Journal of Economic Theory*, vol. 98, pp. 26–54, 2001.
- [95] E. Hazan, “The convex optimization approach to regret minimization,” in *Optimization for Machine Learning*, (S. Sra, S. Nowozin, and S. Wright, eds.), pp. 287–303, MIT Press, 2011.
- [96] E. Hazan and S. Kale, “Better algorithms for benign bandits,” in *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 38–47, 2009.
- [97] E. Hazan and S. Kale, “NEWTRON: An efficient bandit algorithm for online multiclass prediction,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [98] E. Hazan, S. Kale, and M. Warmuth, “Learning rotations with little regret,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2010.
- [99] D. Helmbold and S. Panizza, “Some label efficient learning results,” in *Proceedings of the Annual Conference on Computational Learning Theory (COLT)*, 1997.
- [100] D. P. Helmbold and M. Warmuth, “Learning permutations with exponential weights,” *Journal of Machine Learning Research*, vol. 10, pp. 1705–1736, 2009.
- [101] M. Herbster and M. Warmuth, “Tracking the best expert,” *Machine Learning*, vol. 32, pp. 151–178, 1998.
- [102] J.-B. Hiriart-Urruty and C. Lemaréchal, *Fundamentals of Convex Analysis*. Springer, 2001.
- [103] J. Honda and A. Takemura, “An asymptotically optimal bandit algorithm for bounded support models,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2010.
- [104] J.-B. Hoock and O. Teytaud, “Bandit-based genetic programming,” in *Proceedings of the European Conference on Genetic Programming (EuroGP)*, 2010.
- [105] T. Jaksch, R. Ortner, and P. Auer, “Near-optimal regret bounds for reinforcement learning,” *Journal of Machine Learning Research*, vol. 11, pp. 1563–1600, 2010.

- [106] A. Juditsky, A. Nazin, A. Tsybakov, and N. Vayatis, “Recursive aggregation of estimators by the Mirror Descent algorithm with averaging,” *Problems of Information Transmission*, vol. 41, pp. 368–384, 2005.
- [107] S. Kakade, S. Shalev-Shwartz, and A. Tewari, “Regularization techniques for learning with matrices,” *Journal of Machine Learning Research*, vol. 13, pp. 1865–1890, 2012.
- [108] S. M. Kakade, “On the sample complexity of reinforcement learning,” PhD thesis, Gatsby Computational Neuroscience Unit, University College London, 2003.
- [109] S. M. Kakade, S. Shalev-Shwartz, and A. Tewari, “Efficient bandit algorithms for online multiclass prediction,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2008.
- [110] A. Kalai and S. Vempala, “Efficient algorithms for online decision problems,” *Journal of Computer and System Sciences*, vol. 71, pp. 291–307, 2005.
- [111] S. Kale, L. Reyzin, and R. Schapire, “Non-stochastic bandit slate problems,” in *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [112] V. Kanade, B. McMahan, and B. Bryan, “Sleeping experts and bandits with stochastic action availability and adversarial rewards,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 5, 2009.
- [113] V. Kanade and T. Steinke, “Learning hurdles for sleeping experts,” in *Proceedings of the Innovations in Theoretical Computer Science Conference*, 2012.
- [114] E. Kaufmann, O. Cappé, and A. Garivier, “On Bayesian upper confidence bounds for bandits problems,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 22, 2012.
- [115] E. Kaufmann, N. Korda, and R. Munos, “Thompson sampling: An asymptotically optimal finite-time analysis,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2012.
- [116] M. Kearns, Y. Mansour, and A. Y. Ng, “A sparse sampling algorithm for near-optimal planning in large Markovian decision processes,” *Machine Learning*, vol. 49, pp. 193–208, 2002.
- [117] J. Kiefer, “Sequential minimax search for a maximum,” *Proceedings of the American Mathematical Society*, vol. 4, no. 3, pp. 502–506, 1953.
- [118] J. Kiefer and J. Wolfowitz, “Stochastic estimation of the maximum of a regression function,” *Annals of Mathematical Statistics*, vol. 23, pp. 462–466, 1952.
- [119] J. Kivinen and M. Warmuth, “Relative loss bounds for multidimensional regression problems,” *Machine Learning*, vol. 45, pp. 301–329, 2001.
- [120] R. Kleinberg, “Nearly tight bounds for the continuum-armed bandit problem,” in *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- [121] R. Kleinberg, A. Niculescu-Mizil, and Y. Sharma, “Regret bounds for sleeping experts and bandits,” *Machine Learning*, vol. 80, pp. 245–272, 2010.
- [122] R. Kleinberg, A. Slivkins, and E. Upfal, “Multi-armed bandits in metric spaces,” in *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2008.

- [123] L. Kocsis and C. Szepesvári, “Bandit based Monte-Carlo planning,” in *Proceedings of the European Conference on Machine Learning (ECML)*, 2006.
- [124] W. Koolen, M. Warmuth, and J. Kivinen, “Hedging structured concepts,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2010.
- [125] T. L. Lai and H. Robbins, “Asymptotically efficient adaptive allocation rules,” *Advances in Applied Mathematics*, vol. 6, pp. 4–22, 1985.
- [126] J. Langford and T. Zhang, “The epoch-greedy algorithm for contextual multi-armed bandits,” in *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- [127] P. Lezaud, “Chernoff-type bound for finite Markov chains,” *Annals of Applied Probability*, vol. 8, pp. 849–867, 1998.
- [128] L. Li, W. Chu, J. Langford, and R. E. Schapire, “A contextual-bandit approach to personalized news article recommendation,” in *Proceedings of the International Conference on World Wide Web (WWW)*, 2010.
- [129] K. Liu, Q. Zhao, and B. Krishnamachari, “Dynamic multichannel access with imperfect channel state detection,” *IEEE Transactions on Signal Processing*, vol. 58, pp. 2795–2808, 2010.
- [130] A. Mahajan and D. Teneketzis, “Multi-armed bandit problems,” in *Foundations and Applications of Sensor Management*, pp. 121–151, Springer, 2008.
- [131] O. Maillard and R. Munos, “Adaptive bandits: Towards the best history-dependent strategy,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings Volume 15, 2011.
- [132] O.-A. Maillard, R. Munos, and G. Stoltz, “A finite-time analysis of multi-armed bandits problems with Kullback-Leibler divergences,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 19, 2011.
- [133] S. Mannor and O. Shamir, “From bandits to experts: On the value of side-observations,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [134] S. Mannor and J. N. Tsitsiklis, “The sample complexity of exploration in the multi-armed bandit problem,” *Journal of Machine Learning Research*, vol. 5, pp. 623–648, 2004.
- [135] H. McMahan and A. Blum, “Online geometric optimization in the bandit setting against an adaptive adversary,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2004.
- [136] H. B. McMahan and M. Streeter, “Tighter bounds for multi-armed bandits with expert advice,” in *Proceedings of the International Conference on Learning Theory (COLT)*, 2009.
- [137] A. Nemirovski, “Efficient methods for large-scale convex optimization problems,” *Ekonomika i Matematicheskie Metody*, vol. 15, 1979. (In Russian).
- [138] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, “Robust stochastic approximation approach to stochastic programming,” *SIAM Journal on Optimization*, vol. 19, pp. 1574–1609, 2009.
- [139] A. Nemirovski and D. Yudin, *Problem Complexity and Method Efficiency in Optimization*. Wiley Interscience, 1983.

- [140] Y. Nesterov, “Random gradient-free minimization of convex functions,” Core discussion papers, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 2011.
- [141] G. Neu, A. Gyorgy, C. Szepesvari, and A. Antos, “Online Markov decision processes under bandit feedback,” in *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [142] R. Ortner, D. Ryabko, P. Auer, and R. Munos, “Regret bounds for restless Markov bandits,” arXiv preprint arXiv:1209.2693, 2012.
- [143] V. Perchet and P. Rigollet, “The multi-armed bandit problem with covariates,” Arxiv preprint arXiv:1110.6084, 2011.
- [144] D. Pucci de Farias and N. Megiddo, “Combining expert advice in reactive environments,” *Journal of the ACM*, vol. 53, no. 5, pp. 762–799, 2006.
- [145] A. Rakhlin, “Lecture notes on online learning,” 2009.
- [146] H. Robbins, “Some aspects of the sequential design of experiments,” *Bulletin of the American Mathematics Society*, vol. 58, pp. 527–535, 1952.
- [147] H. Robbins and S. Monro, “A stochastic approximation method,” *Annals of Mathematical Statistics*, vol. 22, pp. 400–407, 1951.
- [148] P. Rusmevichientong and J. Tsitsiklis, “Linearly parameterized bandits,” *Mathematics of Operations Research*, vol. 35, pp. 395–411, 2010.
- [149] A. Salomon and J.-Y. Audibert, “Deviations of stochastic bandit regret,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2011.
- [150] Y. Seldin, P. Auer, F. Laviolette, J. Shawe-Taylor, and R. Ortner, “PAC-Bayesian analysis of contextual bandits,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [151] S. Shalev-Shwartz, “Online learning: Theory, algorithms, and applications,” PhD thesis, The Hebrew University of Jerusalem, 2007.
- [152] A. Slivkins, “Contextual bandits with similarity information,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, JMLR Workshop and Conference Proceedings Volume 19, 2011.
- [153] A. Slivkins and E. Upfal, “Adapting to a changing environment: The Brownian restless bandits,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2008.
- [154] N. Srebro, K. Sridharan, and A. Tewari, “On the universality of online Mirror Descent,” in *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [155] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, “Gaussian process optimization in the bandit setting: No regret and experimental design,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2010.
- [156] G. Stoltz, “Incomplete information and internal regret in prediction of individual sequences,” PhD thesis, Université Paris-Sud, 2005.
- [157] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- [158] C. Szepesvári, “Algorithms for Reinforcement Learning,” Morgan and Claypool, 2010.

- [159] E. Takimoto and M. Warmuth, “Paths kernels and multiplicative updates,” *Journal of Machine Learning Research*, vol. 4, pp. 773–818, 2003.
- [160] C. Tekin and M. Liu, “Online learning of rested and restless bandits,” *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5588–5611, 2012.
- [161] F. Teytaud and O. Teytaud, “Creating an upper-confidence-tree program for Havannah,” in *Advances in Computer Games*, pp. 65–74, 2009.
- [162] W. Thompson, “On the likelihood that one unknown probability exceeds another in view of the evidence of two samples,” *Bulletin of the American Mathematics Society*, vol. 25, pp. 285–294, 1933.
- [163] T. Uchiya, A. Nakamura, and M. Kudo, “Algorithms for adversarial bandit problems with multiple plays,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2010.
- [164] A. Wald, *Sequential Analysis*. J. Wiley and Sons, 1947.
- [165] C. C. Wang, S. R. Kulkarni, and H. V. Poor, “Arbitrary side observations in bandit problems,” *Advances in Applied Mathematics*, vol. 34, no. 4, pp. 903–938, 2005.
- [166] C. C. Wang, S. R. Kulkarni, and H. V. Poor, “Bandit problems with side observations,” *IEEE Transactions on Automatic Control*, vol. 50, no. 3, pp. 338–355, 2005.
- [167] Y. Wang, J.-Y. Audibert, and R. Munos, “Algorithms for infinitely many-armed bandits,” in *Advances in Neural Information Processing Systems (NIPS)*, 2008.
- [168] M. Warmuth, W. Koolen, and D. Helmbold, “Combining initial segments of lists,” in *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2011.
- [169] M. Warmuth and D. Kuzmin, “Randomized online PCA algorithms with regret bounds that are logarithmic in the dimension,” *Journal of Machine Learning Research*, vol. 9, pp. 2287–2320, 2008.
- [170] C. A. Wilkens and B. Sivan, “Single-call mechanisms,” in *ACM Conference on Electronic Commerce (EC)*, 2012.
- [171] J. Y. Yu and S. Mannor, “Unimodal bandits,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2011.
- [172] J. Y. Yu, S. Mannor, and N. Shimkin, “Markov decision processes with arbitrary reward processes,” *Mathematics of Operations Research*, vol. 34, pp. 737–757, 2009.
- [173] Y. Yue, J. Broder, R. Kleinberg, and T. Joachims, “The k -armed dueling bandits problem,” in *Proceedings of the Annual Conference on Learning Theory (COLT)*, 2009.
- [174] Y. Yue and T. Joachims, “Beat the mean bandit,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2011.
- [175] M. Zinkevich, “Online convex programming and generalized infinitesimal gradient Ascent,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2003.