
Online Learning with Switching Costs and Other Adaptive Adversaries

Nicolò Cesa-Bianchi
Università degli Studi di Milano
Italy

Ofer Dekel
Microsoft Research
USA

Ohad Shamir
Microsoft Research
and the Weizmann Institute

Abstract

We study the power of different types of adaptive (nonoblivious) adversaries in the setting of prediction with expert advice, under both full-information and bandit feedback. We measure the player’s performance using a new notion of regret, also known as policy regret, which better captures the adversary’s adaptiveness to the player’s behavior. In a setting where losses are allowed to drift, we characterize—in a nearly complete manner—the power of adaptive adversaries with bounded memories and switching costs. In particular, we show that with switching costs, the attainable rate with bandit feedback is $\tilde{\Theta}(T^{2/3})$. Interestingly, this rate is significantly worse than the $\Theta(\sqrt{T})$ rate attainable with switching costs in the full-information case. Via a novel reduction from experts to bandits, we also show that a bounded memory adversary can force $\tilde{\Theta}(T^{2/3})$ regret even in the full information case, proving that switching costs are easier to control than bounded memory adversaries. Our lower bounds rely on a new stochastic adversary strategy that generates loss processes with strong dependencies.

1 Introduction

An important instance of the framework of prediction with expert advice—see, e.g., [8]—is defined as the following repeated game, between a randomized player with a finite and fixed set of available actions and an adversary. At the beginning of each round of the game, the adversary assigns a loss to each action. Next, the player defines a probability distribution over the actions, draws an action from this distribution, and suffers the loss associated with that action. The player’s goal is to accumulate loss at the smallest possible rate, as the game progresses. Two versions of this game are typically considered: in the *full-information feedback* version, at the end of each round, the player observes the adversary’s assignment of loss values to each action. In the *bandit feedback* version, the player only observes the loss associated with his chosen action, but not the loss values of other actions.

We assume that the adversary is *adaptive* (also called *nonoblivious* by [8] or *reactive* by [16]), which means that the adversary chooses the loss values on round t based on the player’s actions on rounds $1 \dots t - 1$. We also assume that the adversary is deterministic and has unlimited computational power. These assumptions imply that the adversary can specify his entire strategy before the game begins. In other words, the adversary can perform all of the calculations needed to specify, in advance, how he plans to react on each round to any sequence of actions chosen by the player.

More formally, let \mathcal{A} denote the finite set of actions and let X_t denote the player’s random action on round t . We adopt the notation $X_{1:t}$ as shorthand for the sequence $X_1 \dots X_t$. We assume that the adversary defines, in advance, a sequence of *history-dependent loss functions* f_1, f_2, \dots . The input to each loss function f_t is the entire history of the player’s actions so far, therefore the player’s loss on round t is $f_t(X_{1:t})$. Note that the player doesn’t observe the functions f_t , only the losses that result from his past actions. Specifically, in the bandit feedback model, the player observes $f_t(X_{1:t})$ on round t , whereas in the full-information model, the player observes $f_t(X_{1:t-1}, x)$ for all $x \in \mathcal{A}$.

On any round T , we evaluate the player's performance so far using the notion of *regret*, which compares his cumulative loss on the first T rounds to the cumulative loss of the best fixed action in hindsight. Formally, the player's regret on round T is defined as

$$R_T = \sum_{t=1}^T f_t(X_{1:t}) - \min_{x \in \mathcal{A}} \sum_{t=1}^T f_t(x \dots x) . \quad (1)$$

R_T is a random variable, as it depends on the randomized action sequence $X_{1:t}$. Therefore, we also consider the *expected regret* $\mathbb{E}[R_T]$. This definition is the same as the one used in [18] and [3] (in the latter, it is called *policy regret*), but differs from the more common definition of expected regret

$$\mathbb{E} \left[\sum_{t=1}^T f_t(X_{1:t}) - \min_{x \in \mathcal{A}} \sum_{t=1}^T f_t(X_{1:t-1}, x) \right] . \quad (2)$$

The definition in Eq. (2) is more common in the literature (e.g., [4, 17, 10, 16]), but is clearly inadequate for measuring a player's performance against an adaptive adversary. Indeed, if the adversary is adaptive, the quantity $f_t(X_{1:t-1}, x)$ is hardly interpretable —see [3] for a more detailed discussion.

In general, we seek algorithms for which $\mathbb{E}[R_T]$ can be bounded by a sublinear function of T , implying that the per-round expected regret, $\mathbb{E}[R_T]/T$, tends to zero. Unfortunately, [3] shows that arbitrary adaptive adversaries can easily force the regret to grow linearly. Thus, we need to focus on (reasonably) weaker adversaries, which have constraints on the loss functions they can generate.

The weakest adversary we discuss is the *oblivious adversary*, which determines the loss on round t based only on the current action X_t . In other words, this adversary is oblivious to the player's past actions. Formally, the oblivious adversary is constrained to choose a sequence of loss functions that satisfies $\forall t, \forall x_{1:t} \in \mathcal{A}^t$, and $\forall x'_{1:t-1} \in \mathcal{A}^{t-1}$,

$$f_t(x_{1:t}) = f_t(x'_{1:t-1}, x_t) . \quad (3)$$

The majority of previous work in online learning focuses on oblivious adversaries. When dealing with oblivious adversaries, we denote the loss function by ℓ_t and omit the first $t-1$ arguments. With this notation, the loss at time t is simply written as $\ell_t(X_t)$.

For example, imagine an investor that invests in a single stock at a time. On each trading day he invests in one stock and suffers losses accordingly. In this example, the investor is the player and the stock market is the adversary. If the investment amount is small, the investor's actions will have no measurable effect on the market, so the market is oblivious to the investor's actions. Also note that this example relates to the full-information feedback version of the game, as the investor can see the performance of each stock at the end of each trading day.

A stronger adversary is the *oblivious adversary with switching costs*. This adversary is similar to the oblivious adversary defined above, but charges the player an additional switching cost of 1 whenever $X_t \neq X_{t-1}$. More formally, this adversary defines his sequence of loss functions in two steps: first he chooses an oblivious sequence of loss functions, ℓ_1, ℓ_2, \dots , which satisfies the constraint in Eq. (3). Then, he sets $f_1(x) = \ell_1(x)$, and

$$\forall t \geq 2, \quad f_t(x_{1:t}) = \ell_t(x_t) + \mathbb{I}_{\{x_t \neq x_{t-1}\}} . \quad (4)$$

This is a very natural setting. For example, let us consider again the single-stock investor, but now assume that each trade has a fixed commission cost. If the investor keeps his position in a stock for multiple trading days, he is exempt from any additional fees, but when he sells one stock and buys another, he incurs a fixed commission. More generally, this setting (or simple generalizations of it) allows us to capture any situation where choosing a different action involves a costly change of state. In the paper, we will also discuss a special case of this adversary, where the loss function $\ell_t(x)$ for each action is sampled i.i.d. from a fixed distribution.

The switching costs adversary defines f_t to be a function of X_t and X_{t-1} , and is therefore a special case of a more general adversary called an *adaptive adversary with a memory of 1*. This adversary is constrained to choose loss functions that satisfy $\forall t, \forall x_{1:t} \in \mathcal{A}^t$, and $\forall x'_{1:t-2} \in \mathcal{A}^{t-2}$,

$$f_t(x_{1:t}) = f_t(x'_{1:t-2}, x_{t-1}, x_t) . \quad (5)$$

This adversary is more general than the switching costs adversary because his loss functions can depend on the previous action in an arbitrary way. We can further strengthen this adversary and

define the *bounded memory adaptive adversary*, which has a bounded memory of an arbitrary size. In other words, this adversary is allowed to set his loss function based on the player’s m most recent past actions, where m is a predefined parameter. Formally, the bounded memory adversary must choose loss functions that satisfy, $\forall t, \forall x_{1:t} \in \mathcal{A}^t$, and $\forall x'_{1:t-m-1} \in \mathcal{A}^{t-m-1}$,

$$f_t(x_{1:t}) = f_t(x'_{1:t-m-1}, x_{t-m:t}) .$$

In the information theory literature, this setting is called *individual sequence prediction against loss functions with memory* [18].

In addition to the adversary types described above, the bounded memory adaptive adversary has additional interesting special cases. One of them is the *delayed feedback oblivious adversary* of [19], which defines an oblivious loss sequence, but reveals each loss value with a delay of m rounds. Since the loss at time t depends on the player’s action at time $t - m$, this adversary is a special case of a bounded memory adversary with a memory of size m . The delayed feedback adversary is not a focus of our work, and we present it merely as an interesting special case.

So far, we have defined a succession of adversaries of different strengths. This paper’s goal is to understand the upper and lower bounds on the player’s regret when he faces these adversaries. Specifically, we focus on how the expected regret depends on the number of rounds, T , with either full-information or bandit feedback.

1.1 The Current State of the Art

Different aspects of this problem have been previously studied and the known results are surveyed below and summarized in Table 1. Most of these previous results rely on the additional assumption that the range of the loss functions is bounded in a fixed interval, say $[0, C]$. We explicitly make note of this because our new results require us to slightly generalize this assumption.

As mentioned above, the oblivious adversary has been studied extensively and is the best understood of all the adversaries discussed in this paper. With full-information feedback, both the *Hedge* algorithm [15, 11] and the *follow the perturbed leader (FPL)* algorithm [14] guarantee a regret of $\mathcal{O}(\sqrt{T})$, with a matching lower bound of $\Omega(\sqrt{T})$ —see, e.g., [8]. Analyses of Hedge in settings where the loss range may vary over time have also been considered —see, e.g., [9]. The oblivious setting with bandit feedback, where the player only observes the incurred loss $f_t(X_{1:t})$, is called the *nonstochastic (or adversarial) multi-armed bandit* problem. In this setting, the Exp3 algorithm of [4] guarantees the same regret $\mathcal{O}(\sqrt{T})$ as the full-information setting, and clearly the full-information lower bound $\Omega(\sqrt{T})$ still applies.

The *follow the lazy leader (FLL)* algorithm of [14] is designed for the switching costs setting with full-information feedback. The analysis of FLL guarantees that the oblivious component of the player’s expected regret (without counting the switching costs), as well as the expected number of switches, is upper bounded by $\mathcal{O}(\sqrt{T})$, implying an expected regret of $\mathcal{O}(\sqrt{T})$.

The work in [3] focuses on the bounded memory adversary with bandit feedback and guarantees an expected regret of $\mathcal{O}(T^{2/3})$. This bound naturally extends to the full-information setting. We note that [18, 12] study this problem in a different feedback model, which we call *counterfactual feedback*, where the player receives a full description of the history-dependent function f_t at the end of round t . In this setting, the algorithm presented in [12] guarantees an expected regret of $\mathcal{O}(\sqrt{T})$.

Learning with bandit feedback and switching costs has mostly been considered in the economics literature, using a different setting than ours and with prior knowledge assumptions (see [13] for an overview). The setting of stochastic oblivious adversaries (i.e., oblivious loss functions sampled i.i.d. from a fixed distribution) was first studied by [2], where they show that $\mathcal{O}(\log T)$ switches are sufficient to asymptotically guarantee logarithmic regret. The paper [20] achieves logarithmic regret nonasymptotically with $\mathcal{O}(\log T)$ switches.

Several other papers discuss online learning against “adaptive” adversaries [4, 10, 16, 17], but these results are not relevant to our work and can be easily misunderstood. For example, several bandit algorithms have extensions to the “adaptive” adversary case, with a regret upper bound of $\mathcal{O}(\sqrt{T})$ [1]. This bound doesn’t contradict the $\Omega(T)$ lower bound for general adaptive adversaries mentioned

	oblivious	switching cost	memory of size 1	bounded memory	adaptive
Full-Information Feedback					
$\tilde{\mathcal{O}}$	\sqrt{T}	\sqrt{T}	$T^{2/3}$	$T^{2/3}$	T
Ω	\sqrt{T}	\sqrt{T}	\sqrt{T}	$\sqrt{T} \rightarrow \mathbf{T^{2/3}}$	T
Bandit Feedback					
$\tilde{\mathcal{O}}$	\sqrt{T}	$T^{2/3}$	$T^{2/3}$	$T^{2/3}$	T
Ω	\sqrt{T}	$\sqrt{T} \rightarrow \mathbf{T^{2/3}}$	$\sqrt{T} \rightarrow \mathbf{T^{2/3}}$	$\sqrt{T} \rightarrow \mathbf{T^{2/3}}$	T

Table 1: State-of-the-art upper and lower bounds on regret (as a function of T) against different adversary types. Our contribution to this table is presented in bold face.

earlier, since these papers use the regret defined in Eq. (2) rather than the regret used in our work, defined in Eq. (1).

Another related body of work lies in the field of competitive analysis —see [5], which also deals with loss functions that depend on the player’s past actions, and the adversary’s memory may even be unbounded. However, obtaining sublinear regret is generally impossible in this case. Therefore, competitive analysis studies much weaker performance metrics such as the competitive ratio, making it orthogonal to our work.

1.2 Our Contribution

In this paper, we make the following contributions (see Table 1):

- Our main technical contribution is a new lower bound on regret that matches the existing upper bounds in several of the settings discussed above. Specifically, our lower bound applies to the switching costs adversary with bandit feedback and to all strictly stronger adversaries.
- Building on this lower bound, we prove another regret lower bound in the bounded memory setting with *full-information* feedback, again matching the known upper bound.
- We confirm that existing upper bounds on regret hold in our setting and match the lower bounds up to logarithmic factors.
- Despite the lower bound, we show that for switching costs and bandit feedback, if we also assume *stochastic i.i.d. losses*, then one can get a distribution-free regret bound of $\mathcal{O}(\sqrt{T} \log \log \log T)$ for finite action sets, with only $\mathcal{O}(\log \log T)$ switches. This result uses ideas from [7], and is deferred to the supplementary material.

Our new lower bound is a significant step towards a complete understanding of adaptive adversaries; observe that the upper and lower bounds in Table 1 essentially match in all but one of the settings.

Our results have two important consequences. First, observe that the optimal regret against the switching costs adversary is $\Theta(\sqrt{T})$ with full-information feedback, versus $\Theta(T^{2/3})$ with bandit feedback. To the best of our knowledge, this is the first theoretical confirmation that learning with bandit feedback is strictly harder than learning with full-information, even on a small finite action set and even in terms of the dependence on T (previous gaps we are aware of were either in terms of the number of actions [4], or required large or continuous action spaces —see, e.g., [6, 21]). Moreover, recall the regret bound of $\mathcal{O}(\sqrt{T} \log \log \log T)$ against the stochastic i.i.d. adversary with switching costs and bandit feedback. This demonstrates that dependencies in the loss process must play a crucial role in controlling the power of the switching costs adversary. Indeed, the $\Omega(T^{2/3})$ lower bound proven in the next section heavily relies on such dependencies.

Second, observe that in the full-information feedback case, the optimal regret against a switching costs adversary is $\Theta(\sqrt{T})$, whereas the optimal regret against the more general bounded memory adversary is $\Omega(T^{2/3})$. This is somewhat surprising given the ideas presented in [18] and later extended in [3]: The main technique used in these papers is to take an algorithm originally designed for oblivious adversaries, forcefully prevent it from switching actions very often, and obtain a new algorithm that guarantees a regret of $\mathcal{O}(T^{2/3})$ against bounded memory adversaries. This would

seem to imply that a small number of switches is the key to dealing with general bounded memory adversaries. Our result contradicts this intuition by showing that controlling the number of switches is easier than dealing with a general bounded memory adversary.

As noted above, our lower bounds require us to slightly weaken the standard technical assumption that loss values lie in a fixed interval $[0, C]$. We replace it with the following two assumptions:

1. *Bounded range.* We assume that the loss values *on each individual round* are bounded in an interval of constant size C , but we allow this interval to drift from round to round. Formally, $\forall t, \forall x_{1:t} \in \mathcal{A}^t$ and $\forall x'_{1:t} \in \mathcal{A}^t$,

$$|f_t(x_{1:t}) - f_t(x'_{1:t})| \leq C . \quad (6)$$

2. *Bounded drift.* We also assume that the drift of each individual action from round to round is contained in a bounded interval of size D_t , where D_t may grow slowly, as $\mathcal{O}(\sqrt{\log(t)})$. Formally, $\forall t$ and $\forall x_{1:t} \in \mathcal{A}^t$,

$$|f_t(x_{1:t}) - f_{t+1}(x_{1:t}, x_t)| \leq D_t . \quad (7)$$

Since these assumptions are a relaxation of the standard assumption, all of the known lower bounds on regret automatically extend to our relaxed setting. For our results to be consistent with the current state of the art, we must also prove that all of the known upper bounds continue to hold after the relaxation, up to logarithmic factors.

2 Lower Bounds

In this section, we prove lower bounds on the player's expected regret in various settings.

2.1 $\Omega(T^{2/3})$ with Switching Costs and Bandit Feedback

We begin with a $\Omega(T^{2/3})$ regret lower bound against an oblivious adversary with switching costs, when the player receives bandit feedback. It is enough to consider a very simple setting, with only two actions, labeled 1 and 2. Using the notation introduced earlier, we use ℓ_1, ℓ_2, \dots to denote the oblivious sequence of loss functions chosen by the adversary before adding the switching cost.

Theorem 1. *For any player strategy that relies on bandit feedback and for any number of rounds T , there exist loss functions f_1, \dots, f_T that are oblivious with switching costs, with a range bounded by $C = 2$, and a drift bounded by $D_t = \sqrt{3 \log(t) + 16}$, such that $\mathbb{E}[R_T] \geq \frac{1}{40} T^{2/3}$.*

The full proof is given in the supplementary material, and here we give an informal proof sketch. We begin by constructing a *randomized* adversarial strategy, where the loss functions ℓ_1, \dots, ℓ_T are an instantiation of random variables L_t, \dots, L_T defined as follows. Let ξ_1, \dots, ξ_T be i.i.d. standard Gaussian random variables (with zero mean and unit variance) and let Z be a random variable that equals -1 or 1 with equal probability. Using these random variables, define for all $t = 1 \dots T$

$$L_t(1) = \sum_{s=1}^t \xi_s , \quad L_t(2) = L_t(1) + ZT^{-1/3} . \quad (8)$$

In words, $\{L_t(1)\}_{t=1}^T$ is simply a Gaussian random walk and $\{L_t(2)\}_{t=1}^T$ is the same random walk, slightly shifted up or down—see figure 1 for an illustration. It is straightforward to confirm that this loss sequence has a bounded range, as required by the theorem: by construction we have $|\ell_t(1) - \ell_t(2)| = T^{-1/3} \leq 1$ for all t , and since the switching cost can add at most 1 to the loss on each round, we conclude that $|f_t(1) - f_t(2)| \leq 2$ for all t . Next, we show that the expected regret of any player against this random loss sequence is $\Omega(T^{2/3})$, where expectation is taken over the randomization of both the adversary and the player. The intuition is that the player can only gain information about which action is better by switching between them. Otherwise, if he stays on the same action, he only observes a random walk, and gets no further information. Since the gap between the two losses on each round is $T^{-1/3}$, the player must perform $\Omega(T^{2/3})$ switches before he can identify the better action. If the player performs that many switches, the total regret incurred due to the switching costs is $\Omega(T^{2/3})$. Alternatively, if the player performs $o(T^{2/3})$ switches, he




Figure 1: A particular realization of the random loss sequence defined in Eq. (8). The sequence of losses for action 1 follows a Gaussian random walk, whereas the sequence of losses for action 2 follows the same random walk, but slightly shifted either up or down.

can't identify the better action; as a result he suffers an expected regret of $\Omega(T^{-1/3})$ on each round and a total regret of $\Omega(T^{2/3})$.

Since the randomized loss sequence defined in Eq. (8), plus a switching cost, achieves an expected regret of $\Omega(T^{2/3})$, there must exist at least one *deterministic* loss sequence $\ell_1 \dots \ell_T$ with a regret of $\Omega(T^{2/3})$. In our proof, we show that there exists such $\ell_1 \dots \ell_T$ with bounded drift.

2.2 $\Omega(T^{2/3})$ with Bounded Memory and Full-Information Feedback

We build on Thm. 1 and prove a $\Omega(T^{2/3})$ regret lower bound in the *full-information setting*, where we get to see the entire loss vector on every round. To get this strong result, we need to give the adversary a little bit of extra power: memory of size 2 instead of size 1 as in the case of switching costs. To show this result, we again consider a simple setting with two actions.

Theorem 2. *For any player strategy that relies on full-information feedback and for any number of rounds $T \geq 2$, there exist loss functions f_1, \dots, f_T , each with a memory of size $m = 2$, a range bounded by $C = 2$, and a drift bounded by $D_t = \sqrt{3 \log(t) + 18}$, such that $\mathbb{E}[R_T] \geq \frac{1}{40}(T-1)^{2/3}$.*

The formal proof is deferred to the supplementary material and a proof sketch is given here. The proof is based on a reduction from full-information to bandit feedback that might be of independent interest. We construct the adversarial loss sequence as follows: on each round, the adversary assigns the *same* loss to both actions. Namely, the value of the loss depends only on the player's previous two actions, and not on his action on the current round. Recall that even in the full-information version of the game, the player doesn't know what the losses would have been had he chosen different actions in the past. Therefore, we have made the full-information game as difficult as the bandit game. Specifically, we construct an oblivious loss sequence $\ell_1 \dots \ell_T$ as in Thm. 1 and define

$$f_t(x_{1:t}) = \ell_{t-1}(x_{t-1}) + \mathbb{I}_{\{x_{t-1} \neq x_{t-2}\}}. \quad (9)$$

In words, we define the loss on round t of the full-information game to be equal to the loss on round $t-1$ of a bandits-with-switching-costs game in which the player chooses the same sequence of actions. This can be done with a memory of size 2, since the loss in Eq. (9) is fully specified by the player's choices on rounds $t, t-1, t-2$. Therefore, the $\Omega(T^{2/3})$ lower bound for switching costs and bandit feedback extends to the full-information setting with a memory of size at least 2.

3 Upper Bounds

In this section, we show that the known upper bounds on regret, originally proved for bounded losses, can be extended to the case of losses with bounded range and bounded drift. Specifically, of the upper bounds that appear in Table 1, we prove the following:

- $\mathcal{O}(\sqrt{T})$ for an oblivious adversary with switching costs, with full-information feedback.
- $\tilde{\mathcal{O}}(\sqrt{T})$ for an oblivious adversary with bandit feedback (where $\tilde{\mathcal{O}}$ hides logarithmic factors).
- $\tilde{\mathcal{O}}(T^{2/3})$ for a bounded memory adversary with bandit feedback.

The remaining upper bounds in Table 1 are either trivial or follow from the principle that an upper bound still holds if we weaken the adversary or provide a more informative feedback.

3.1 $\mathcal{O}(\sqrt{T})$ with Switching Costs and Full-Information Feedback

In this setting, $f_t(x_{1:t}) = \ell_t(x_t) + \mathbb{I}_{\{x_t \neq x_{t-1}\}}$. If the oblivious losses $\ell_1 \dots \ell_T$ (without the additional switching costs) were all bounded in $[0, 1]$, the *Follow the Lazy Leader* (FLL) algorithm of [14] would guarantee a regret of $\mathcal{O}(\sqrt{T})$ with respect to these losses (again, without the additional switching costs). Additionally, FLL guarantees that its expected number of switches is $\mathcal{O}(\sqrt{T})$. We use a simple reduction to extend these guarantees to loss functions with a range bounded in an interval of size C and with an arbitrary drift.

On round t , after choosing an action and receiving the loss function ℓ_t , the player defines the modified loss $\ell'_t(x) = \frac{1}{C-1}(\ell_t(x) - \min_y \ell_t(y))$ and feeds it to the FLL algorithm. The FLL algorithm then chooses the next action.

Theorem 3. *If each of the loss functions f_1, f_2, \dots is oblivious with switching costs and has a range bounded by C then the player strategy described above attains $\mathcal{O}(C\sqrt{T})$ expected regret.*

The full proof is given in the supplementary material but the proof technique is straightforward. We first show that each ℓ'_t is bounded in $[0, 1]$ and therefore the standard regret bound for FLL holds with respect to the sequence of modified loss functions ℓ'_1, ℓ'_2, \dots . Then we show that the guarantees provided for FLL imply a regret of $\mathcal{O}(\sqrt{T})$ with respect to the original loss sequence f_1, f_2, \dots .

3.2 $\tilde{\mathcal{O}}(\sqrt{T})$ with an Oblivious Adversary and Bandit Feedback

In this setting, $f_t(x_{1:t})$ simply equals $\ell_t(x_t)$. The reduction described in the previous subsection cannot be used in the bandit setting, since $\min_x \ell_t(x)$ is unknown to the player, and a different reduction is needed. The player sets a fixed horizon T and focuses on controlling his regret at time T ; he can then use a standard doubling trick [8] to handle an infinite horizon. The player uses the fact that each f_t has a range bounded by C . Additionally, he defines $D = \max_{t \leq T} D_t$ and on each round he defines the modified loss

$$f'_t(x_{1:t}) = \frac{1}{2(C+D)}(\ell_t(x_t) - \ell_{t-1}(x_{t-1})) + \frac{1}{2}. \quad (10)$$

Note that $f'_t(X_{1:t})$ can be computed by the player using only bandit feedback. The player then feeds $f'_t(X_{1:t})$ to an algorithm that guarantees a $\mathcal{O}(\sqrt{T})$ standard regret (see definition in Eq. (2)) against a fixed action. The Exp3 algorithm, due to [4], is such an algorithm. The player chooses his actions according to the choices made by Exp3. The following theorem states that this reduction results in a bandit algorithm that guarantees a regret of $\tilde{\mathcal{O}}(\sqrt{T})$ against oblivious adversaries.

Theorem 4. *If each of the loss functions $f_1 \dots f_T$ is oblivious with a range bounded by C and a drift bounded by $D_t = \mathcal{O}(\sqrt{\log(t)})$ then the player strategy described above attains $\tilde{\mathcal{O}}(C\sqrt{T})$ expected regret.*

The full proof is given in the supplementary material. In a nutshell, we show that each f'_t is a loss function bounded in $[0, 1]$ and that the analysis of Exp3 guarantees a regret of $\mathcal{O}(\sqrt{T})$ with respect to the loss sequence $f'_1 \dots f'_T$. Then, we show that this guarantee implies a regret of $(C+D)\mathcal{O}(\sqrt{T}) = \tilde{\mathcal{O}}(C\sqrt{T})$ with respect to the original loss sequence $f_1 \dots f_T$.

3.3 $\tilde{\mathcal{O}}(T^{2/3})$ with Bounded Memory and Bandit Feedback

Proving an upper bound against an adversary with a memory of size m , with bandit feedback, requires a more delicate reduction. As in the previous section, we assume a finite horizon T and we let $D = \max_t D_t$. Let $K = |\mathcal{A}|$ be the number of actions available to the player.

Since $f_T(x_{1:t})$ depends only on the last $m+1$ actions in $x_{1:t}$, we slightly overload our notation and define $f_t(x_{t-m:t})$ to mean the same as $f_t(x_{1:t})$. To define the reduction, the player fixes a base

action $x_0 \in \mathcal{A}$ and for each $t > m$ he defines the loss function

$$\hat{f}_t(x_{t-m:t}) = \frac{1}{2(C + (m+1)D)} (f_t(x_{t-m:t}) - f_{t-m-1}(x_0 \dots x_0)) + \frac{1}{2}.$$

Next, he divides the T rounds into J consecutive epochs of equal length, where $J = \Theta(T^{2/3})$. We assume that the epoch length T/J is at least $2K(m+1)$, which is true when T is sufficiently large. At the beginning of each epoch, the player plans his action sequence for the entire epoch. He uses some of the rounds in the epoch for exploration and the rest for exploitation. For each action in \mathcal{A} , the player chooses an exploration interval of $2(m+1)$ consecutive rounds within the epoch. These K intervals are chosen randomly, but they are not allowed to overlap, giving a total of $2K(m+1)$ exploration rounds in the epoch. The details of how these intervals are drawn appears in our analysis, in the supplementary material. The remaining $T/J - 2K(m+1)$ rounds are used for exploitation.

The player runs the Hedge algorithm [11] in the background, invoking it only at the beginning of each epoch and using it to choose one exploitation action that will be played consistently on all of the exploitation rounds in the epoch. In the exploration interval for action x , the player first plays $m+1$ rounds of the base action x_0 followed by $m+1$ rounds of the action x . Letting t_x denote the first round in this interval, the player uses the observed losses $f_{t_x+m}(x_0 \dots x_0)$ and $f_{t_x+2m+1}(x \dots x)$ to compute $\hat{f}_{t_x+2m+1}(x \dots x)$. In our analysis, we show that the latter is an unbiased estimate of the average value of $\hat{f}_t(x \dots x)$ over t in the epoch. At the end of the epoch, the K estimates are fed as feedback to the Hedge algorithm.

We prove the following regret bound, with the proof deferred to the supplementary material.

Theorem 5. *If each of the loss functions $f_1 \dots f_T$ has a memory of size m , a range bounded by C , and a drift bounded by $D_t = \mathcal{O}(\sqrt{\log(t)})$ then the player strategy described above attains $\tilde{\mathcal{O}}(T^{2/3})$ expected regret.*

4 Discussion

In this paper, we studied the problem of prediction with expert advice against different types of adversaries, ranging from the oblivious adversary to the general adaptive adversary. We proved upper and lower bounds on the player's regret against each of these adversary types, in both the full-information and the bandit feedback models. Our lower bounds essentially matched our upper bounds in all but one case: the adaptive adversary with a unit memory in the full-information setting, where we only know that regret is $\Omega(\sqrt{T})$ and $\mathcal{O}(T^{2/3})$. Our bounds have two important consequences. First, we characterize the regret attainable with switching costs, and show a setting where predicting with bandit feedback is strictly more difficult than predicting with full-information feedback—even in terms of the dependence on T , and even on small finite action sets. Second, in the full-information setting, we show that predicting against a switching costs adversary is strictly easier than predicting against an arbitrary adversary with a bounded memory. To obtain our results, we had to relax the standard assumption that loss values are bounded in $[0, 1]$. Re-introducing this assumption and proving similar lower bounds remains an elusive open problem. Many other questions remain unanswered. Can we characterize the dependence of the regret on the number of actions? Can we prove regret bounds that hold with high probability? Can our results be generalized to more sophisticated notions of regret, as in [3]?

In addition to the adversaries discussed in this paper, there are other interesting classes of adversaries that lie between the oblivious and the adaptive. A notable example is the family of *deterministically adaptive* adversaries, which includes adversaries that adapt to the player's actions in a known deterministic way, rather than in a secret malicious way. For example, imagine playing a multi-armed bandit game where the loss values are initially oblivious, but whenever the player chooses an arm with zero loss, the loss of the same arm on the next round is deterministically changed to zero. Many real-world online prediction scenarios are deterministically adaptive, but we lack a characterization of the expected regret in this setting.

Acknowledgments

Part of this work was done while NCB was visiting OD at Microsoft Research, whose support is gratefully acknowledged.

References

- [1] J. Abernethy and A. Rakhlin. Beating the adaptive bandit with high probability. In *COLT*, 2009.
- [2] R. Agrawal, M.V. Hedge, and D. Teneketzis. Asymptotically efficient adaptive allocation rules for the multiarmed bandit problem with switching cost. *IEEE Transactions on Automatic Control*, 33(10):899–906, 1988.
- [3] R. Arora, O. Dekel, and A. Tewari. Online bandit learning against an adaptive adversary: from regret to policy regret. In *Proceedings of the Twenty-Ninth International Conference on Machine Learning*, 2012.
- [4] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [5] A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.
- [6] S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12:1655–1695, 2011.
- [7] N. Cesa-Bianchi, C. Gentile, and Y. Mansour. Regret minimization for reserve prices in second-price auctions. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA13)*, 2013.
- [8] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [9] N. Cesa-Bianchi, Y. Mansour, and G. Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2/3):321–352, 2007.
- [10] V. Dani and T. P. Hayes. Robbing the bandit: Less regret in online geometric optimization against an adaptive adversary. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.
- [11] Y. Freund and R.E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and System Sciences*, 55(1):119–139, 1997.
- [12] A. Gyorgy and G. Neu. Near-optimal rates for limited-delay universal lossy source coding. In *IEEE International Symposium on Information Theory*, pages 2218–2222, 2011.
- [13] T. Jun. A survey on the bandit problem with switching costs. *De Economist*, 152:513–541, 2004.
- [14] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71:291–307, 2005.
- [15] N. Littlestone and M.K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.
- [16] O. Maillard and R. Munos. Adaptive bandits: Towards the best history-dependent strategy. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 2010.
- [17] H. B. McMahan and A. Blum. Online geometric optimization in the bandit setting against an adaptive adversary. In *Proceedings of the Seventeenth Annual Conference on Learning Theory*, 2004.
- [18] N. Merhav, E. Ordentlich, G. Seroussi, and M.J. Weinberger. Sequential strategies for loss functions with memory. *IEEE Transactions on Information Theory*, 48(7):1947–1958, 2002.
- [19] C. Mesterharm. Online learning with delayed label feedback. In *Proceedings of the Sixteenth International Conference on Algorithmic Learning Theory*, 2005.
- [20] R. Ortner. Online regret bounds for Markov decision processes with deterministic transitions. *Theoretical Computer Science*, 411(29–30):2684–2695, 2010.
- [21] O. Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. *CoRR*, abs/1209.2388, 2012.

Appendices

A Distribution-free regret bound for bandits with switching costs

In this appendix we adapt results of [7] to show a strategy that achieves $\mathcal{O}(\sqrt{T \log \log \log T})$ regret against any i.i.d. oblivious adversary in the bandit setting with switching costs, assuming a finite action set $\mathcal{A} = \{1 \dots K\}$. The strategy used by this stochastic adversary is specified by a probability distribution over oblivious loss functions. The oblivious loss function for each step $t = 1, 2, \dots$ is the realization on an independent draw L_t from this distribution. The regret of a player choosing actions $X_0 = X_1, X_2, \dots$ is defined by

$$R_T = \sum_{t=1}^T \mathbb{E}_t [L_t(X_t) + \mathbb{I}_{\{X_t \neq X_{t-1}\}}] - \min_{x \in \mathcal{A}} \sum_{t=1}^T \mathbb{E} [L_t(x)]$$

where the expectation \mathbb{E} is over the random draw of each L_t and the possible randomization of the player, and the expectation \mathbb{E}_t is conditioned over $X_1, L_1(X_1), \dots, X_{t-1}, L_{t-1}(X_{t-1})$.

Our result focuses on loss distributions such that the law of each marginal $L_1(x)$ is subgaussian. A random variable Z is subgaussian if there exist constants b, c such that for any $a > 0$ $\mathbb{P}(Z > \mathbb{E} Z + a) \leq be^{-ca^2}$ and $\mathbb{P}(Z < \mathbb{E} Z - a) \leq be^{-ca^2}$. One can then show that, for any i.i.d. sequence Z_1, \dots, Z_T of subgaussian random variables,

$$\mathbb{P} \left(\left| \frac{1}{T} \sum_{t=1}^T Z_t - \mathbb{E} Z_1 \right| > \sqrt{\frac{112b}{cT} \ln \frac{1}{\delta}} \right) \leq \delta. \quad (11)$$

In the following, we use the notation $\mathbb{E}[L_t(x)] = \mu(x)$ and $\mu^* = \min_{x \in \mathcal{A}} \mu(x)$.

Theorem 6. *Consider a finite action set $\mathcal{A} = \{1 \dots K\}$. Then for each T there exists a deterministic player strategy for the bandit game with i.i.d. oblivious adversaries and switching costs, whose regret after T steps is $\mathcal{O}(\sqrt{T \log \log \log T})$ with high probability, provided the distribution of $L_1(x)$ is subgaussian for each $x \in \mathcal{A}$.*

Proof. Consider the following player that proceeds in stages. At each stage $s = 1, 2, \dots, S$, the player maintains a set $A_s \subseteq \mathcal{A}$ of active actions. Each action is played $T_s/|A_s|$ times in a round-robin fashion, where $T_s = T^{1-2^{-s}}$ is the total number of plays in stage s and T is the known horizon. Note that the overall number of switches is at most KS , where

$$S = \min \left\{ j \in \mathbb{N} : \sum_{s=1}^j T_s \geq T \right\} = \mathcal{O}(\ln \ln T).$$

Let $\hat{\mu}_s(x)$ the sample mean of losses for action x in stage s , and define

$$\hat{x}_s = \operatorname{argmin}_{x \in A_s} \hat{\mu}_s(x)$$

the best empirical action in stage s . The sets A_s of active actions are defined as follows: $A_1 = \mathcal{A}$ and

$$A_s = \left\{ x \in A_{s-1} : \hat{\mu}_{s-1}(x) \leq \hat{\mu}_{s-1}(\hat{x}_{s-1}) + 2C_{s-1} \right\}$$

where

$$C_s = \sqrt{112(b/c) \frac{K}{T_s} \ln \frac{KS}{\delta}}.$$

Note that $A_S \subseteq \dots \subseteq A_1$ by construction. Also, using (11) and the union bound we have that

$$\max_{x \in A_s} |\hat{\mu}_s(x) - \mu(x)| \leq C_s \quad (12)$$

simultaneously for all $s = 1, \dots, S$ with probability at least $1 - \delta$.

We claim the following.

Claim 1. *With probability at least $1 - \delta$,*

$$x^* \in \bigcap_{s=1}^S A_s \quad \text{and} \quad 0 \leq \hat{\mu}_s(x^*) - \hat{\mu}_s(\hat{x}_s) \leq 2C_s \quad \text{for all } s = 1, \dots, S.$$

Proof of Claim. We prove the lemma by induction on $s = 1, \dots, S$. We first show that the base case $s = 1$ holds with probability at least $1 - \delta/S$. Then we show that if the claim holds for $s - 1$, then it holds for s with probability at least $1 - \delta/S$ over all random events in stage s . Therefore, using a union bound over $s = 1, \dots, S$ we get that the claim holds simultaneously for all s with probability at least $1 - \delta$.

For the base case $s = 1$ note that $x^* \in A_1$ by definition, and thus $\hat{\mu}_1(\hat{x}_1) \leq \hat{\mu}_1(x^*)$ holds. Moreover, using (12) we obtain that

$$\hat{\mu}_1(x^*) - \mu(x^*) \leq C_1 \quad \text{and} \quad \mu(\hat{x}_1) - \hat{\mu}_1(\hat{x}_1) \leq C_1$$

holds with probability at least $1 - \delta/S$. Since $\mu(x^*) - \mu(\hat{x}_1) \leq 0$ by definition of x^* , we obtain

$$0 \leq \hat{\mu}_1(x^*) - \hat{\mu}_1(\hat{x}_1) \leq 2C_1$$

as required. We now prove the claim for $s > 1$ using the inductive assumption

$$x^* \in A_{s-1} \quad \text{and} \quad 0 \leq \hat{\mu}_{s-1}(x^*) - \hat{\mu}_{s-1}(\hat{x}_{s-1}) \leq 2C_{s-1}.$$

The inductive assumption directly implies that $x^* \in A_s$. Thus we have $\hat{\mu}_i(\hat{x}_s) \leq \hat{\mu}_s(x^*)$, because \hat{x}_s minimizes $\hat{\mu}_s$ over a set that contains x^* . The rest of the proof of the claim closely follows that of the base case $s = 1$. \square

Now, for any $s = 1, \dots, S$ and for any $x \in A_s$ we have that

$$\begin{aligned} \mu(x) - \mu(x^*) &\leq \hat{\mu}_{s-1}(x) - \mu(x^*) + C_{s-1} \quad \text{by (12)} \\ &\leq \hat{\mu}_{s-1}(\hat{x}_{s-1}) - \mu(x^*) + 3C_{s-1} \quad \text{by definition of } A_{s-1}, \text{ since } x \in A_s \subseteq A_{s-1} \\ &\leq \hat{\mu}_{s-1}(x^*) - \mu(x^*) + 3C_{s-1} \quad \text{since } \hat{x}_{s-1} \text{ minimizes } \hat{\mu}_{s-1} \text{ in } A_{s-1} \\ &\leq 4C_{s-1} \quad \text{by (12)} \end{aligned}$$

holds with probability at least $1 - \delta/S$. Hence, recalling that

$$\sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}} \leq KS$$

holds deterministically, the regret of the player over the T plays can be bounded as follows

$$\begin{aligned} KS + \sum_{t=1}^T (\mu(X_t) - \mu^*) &= KS + \sum_{s=1}^S \frac{T_s}{|A_s|} \sum_{x \in A_s} (\mu(x) - \mu^*) \\ &= KS + \frac{T_1}{K} \sum_{i=1}^K (\mu(x) - \mu^*) + \sum_{s=2}^S \frac{T_s}{|A_s|} \sum_{x \in A_s} (\mu(x) - \mu^*) \\ &\leq KS + T_1 \mu^* + \sum_{i=2}^S 4T_s \sqrt{112(b/c) \frac{K}{T_s} \ln \frac{KS}{\delta}} \\ &= KS + T_1 \mu^* + 4 \sqrt{112(b/c) K \ln \frac{KS}{\delta}} \sum_{s=2}^S \frac{T_s}{\sqrt{T_{s-1}}} \end{aligned}$$

Now, since $T_1 = \sqrt{T}$, $T_s/\sqrt{T_{s-1}} = \sqrt{T}$ and $S = \mathcal{O}(\ln \ln T)$, we obtain that with probability at least $1 - \delta$ the regret is at most of order

$$K \ln \ln T + \mu^* \sqrt{T} + \sqrt{KT \left(\ln \frac{K}{\delta} + \ln \ln T \right)}$$

as desired. \square

B Proof of Thm. 1

As mentioned in the text, we first consider the player's expected regret against a randomized adversary. Specifically, we define

$$\forall t \quad L_t(1) = \sum_{s=1}^t \xi_s \quad \text{and} \quad L_t(2) = L_t(1) + Z\epsilon,$$

where $\xi_1 \dots \xi_T$ are independent standard Gaussians, Z equals -1 or 1 with equal probability, and ϵ is the gap between the losses of the two actions (which will later be set to $\epsilon = T^{-1/3}$).

Next, we assume for now, without loss of generality, that the player is deterministic. A deterministic player chooses each action X_t as a deterministic function of the random losses suffered on the previous rounds, $L_1(X_1) \dots L_{t-1}(X_{t-1})$. We can make this assumption because any randomized player strategy can be seen as a distribution over deterministic player strategies, and since the randomization used by the adversary is independent of the player's strategy.

In the results below, \mathbb{P} denotes the distribution of the randomized adversary. We also introduce the conditional distributions $\mathbb{S} = \mathbb{P}(\cdot \mid Z > 0)$ (i.e., 1 is the better action) and $\mathbb{Q} = \mathbb{P}(\cdot \mid Z < 0)$ (i.e., 2 is the better action). Since Z has an equal probability of being negative or positive, it holds that $\mathbb{P} = \frac{1}{2}(\mathbb{S} + \mathbb{Q})$.

We begin with the following technical lemma.

Lemma 1. *Let $\mathbb{I}_{\{x_{t-1} \neq x_t\}}$ indicate whether the player switched actions on round t (and 1 for $t = 1$). Then for any event A ,*

$$|\mathbb{S}(A) - \mathbb{Q}(A)| \leq \epsilon \sqrt{\mathbb{E} \left[\sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}} \right]}$$

where the expectation in the right-hand side is with respect to \mathbb{P} .

Proof. To show this, we use the chain rule for relative entropy, which implies

$$D_{\text{KL}}(\mathbb{S} \parallel \mathbb{Q}) = \sum_{t=1}^T D_{\text{KL}}(\mathbb{S}_{t-1} \parallel \mathbb{Q}_{t-1}) \quad (13)$$

where \mathbb{S}_{t-1} and \mathbb{Q}_{t-1} denote the distributions of the player's loss $L_t(x_t)$ conditioned on L_1, \dots, L_{t-1} , when the joint distribution of L_1, \dots, L_T is, respectively, \mathbb{S} and \mathbb{Q} .

Let us focus on a particular term $D_{\text{KL}}(\mathbb{S}_{t-1} \parallel \mathbb{Q}_{t-1})$ and a particular realization of the random losses L_1, \dots, L_{t-1} . Since we assume a deterministic player strategy, for any such realization the player's choices $x_{1:t}$ are all determined, and we deterministically have that the player either switched or not at time t . If he did not switch, then $L_t(x_t)$ is distributed as $L_{t-1}(x_{t-1}) + \xi_t$ under both measures \mathbb{S}_{t-1} and \mathbb{Q}_{t-1} , so the relative entropy between them is zero. If he did switch, then $L_t(x_t)$ is distributed as $L_{t-1}(x_{t-1}) - \epsilon + \xi$ under \mathbb{S}_{t-1} (where the switch is towards the best action), and as $L_{t-1}(x_{t-1}) + \epsilon + \xi$ under \mathbb{Q}_{t-1} (where the switch is towards the worst action). Hence, the relative entropy is the same as two standard Gaussians whose means are shifted by 2ϵ , namely $2\epsilon^2$. So overall, we can upper bound Eq. (13) by

$$2\epsilon^2 \mathbb{E} \left[\sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}} \mid Z > 0 \right]. \quad (14)$$

Using a similar argument, we also show that $D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{S})$ is upper bounded by Eq. (14) in which the conditioning on $Z > 0$ is replaced by $Z < 0$. Then, Pinsker's inequality implies that $|\mathbb{S}(A) - \mathbb{Q}(A)|^2$ is at most

$$\frac{\epsilon^2}{2} \left(\mathbb{E} \left[\sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}} \mid Z > 0 \right] + \mathbb{E} \left[\sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}} \mid Z < 0 \right] \right) = \epsilon^2 \mathbb{E} \left[\sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}} \right]$$

which gives the desired bound. \square

With this lemma, we can prove a lower bound on the expected regret for randomized adversaries.

Lemma 2. *By picking $\epsilon = T^{-1/3}$, the expected regret of any deterministic player strategy, over the randomness of the adversary, is at least $\frac{1}{10}T^{2/3}$.*

Proof. Let A be the event that the worst action (action 2 if $Z > 0$, and 1 if $Z < 0$) was picked by the player at least $T/2$ times. Also, let $S_T = \sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}}$ be the number of switches the player performs. Then

$$\mathbb{E}[R_T] \geq \mathbb{E} \left[\max \left\{ S_T, \frac{\epsilon T}{2} \mathbb{I}_{\{A\}} \right\} \right] \geq \mathbb{E} \left[\frac{1}{2} \left(S_T + \frac{\epsilon T}{2} \mathbb{I}_{\{A\}} \right) \right] = \frac{1}{2} \mathbb{E}[S_T] + \frac{\epsilon T}{4} \mathbb{P}(A).$$

Moreover, letting A_1 denote the event that the player chose action 1 at least $T/2$ times, and letting A_2 denote the event that the player chose action 2 at least $T/2$ times, we have $\mathbb{P}(A) = \frac{1}{2}(\mathbb{S}(A_2) + \mathbb{Q}(A_1))$. Substituting this, we get

$$\frac{1}{2} \mathbb{E}[S_T] + \frac{\epsilon T}{8} (\mathbb{S}(A_2) + \mathbb{Q}(A_1)).$$

Using Lemma 1 to lower bound $\mathbb{Q}(A_1)$ via $\mathbb{S}(A_1)$, we get a lower bound of

$$\begin{aligned} \frac{1}{2} \mathbb{E}[S_T] + \frac{\epsilon T}{8} (\mathbb{S}(A_2) + \mathbb{S}(A_1) - \epsilon \sqrt{\mathbb{E}[S_T]}) &\geq \frac{1}{2} \mathbb{E}[S_T] + \frac{\epsilon T}{8} (\mathbb{S}(A_1 \cup A_2) - \epsilon \sqrt{\mathbb{E}[S_T]}) \\ &= \frac{1}{2} \mathbb{E}[S_T] + \frac{\epsilon T}{8} (1 - \epsilon \sqrt{\mathbb{E}[S_T]}) = \frac{1}{2} \mathbb{E}[S_T] - \frac{\epsilon^2 T}{8} \sqrt{\mathbb{E}[S_T]} + \frac{\epsilon T}{8}, \end{aligned}$$

where we used a union bound and the fact that either A_1 or A_2 always holds. This is a quadratic function of $\sqrt{\mathbb{E}[S_T]}$, and it is easily verified that the lowest possible value it can attain (for any value of $\mathbb{E}[S_T]$) is

$$\frac{\epsilon T}{8} - \frac{\epsilon^4 T^2}{128}.$$

Picking $\epsilon = T^{-1/3}$, this equals $(\frac{1}{8} - \frac{1}{128}) T^{2/3} > \frac{1}{10} T^{2/3}$. \square

The lemma above tells us that for the randomized adversary strategy we have devised, the expected regret for any deterministic player is at least $\frac{1}{10}T^{2/3}$. This implies that there exist some *deterministic* adversarial strategy, for which the expected regret of any *possibly randomized* player is at least $\frac{1}{10}T^{2/3}$. However, we are not done yet, since this strategy doesn't guarantee that the losses have bounded drift: In our case, the variation is governed by a potentially unbounded Gaussian random variable, so the deterministic adversary strategy that we picked might have an arbitrarily large drift. So now, our goal will be to show that there exists some deterministic adversarial strategy for which the expected regret is large, *and the variation is bounded*. To do this, the plan is to show that the probabilities (over the adversary's strategy) of the two events are large, summing to a number larger than one. This means there is some realization of the losses such that both events occur. We first state and prove two auxiliary lemmas, and then provide two more fundamental lemmas which together give us the required result.

Lemma 3. *Let Y be a random variable in $[-b, b]$ (where $b > 0$), and $\mathbb{E}[Y] \geq c$ for some $c \in [0, b/2]$. Then we have*

$$\mathbb{P}(Y \geq c/2) \geq \frac{c}{2b - c} \geq \frac{c}{2b}.$$

Proof.

$$\begin{aligned} c &\leq \mathbb{E}[Y] = \mathbb{P}(Y \geq c/2) \mathbb{E}[Y | Y \geq c/2] + \mathbb{P}(Y < c/2) \mathbb{E}[Y | Y < c/2] \\ &\leq \mathbb{P}(Y \geq c/2) b + (1 - \mathbb{P}(Y \geq c/2)) c/2 \end{aligned}$$

Solving for $\mathbb{P}(Y \geq c/2)$ gives the desired result. \square

Lemma 4. *Let ξ_1, ξ_2, \dots be an infinite sequence of independent standard Gaussian random variables. Then for any $\delta \in (0, 1)$*

$$\mathbb{P} \left(\exists t : |\xi_t| \geq \sqrt{3 \log(2t/\delta)} \right) \leq \delta.$$

Proof. By a standard Gaussian tail bound, we have that $\mathbb{P}(|\xi_t| > x) \leq \exp(-x^2/2)$ for any $x \geq 0$. This implies that

$$\mathbb{P}(|\xi_t| \geq \sqrt{3 \log(2t/\delta)}) \leq \left(\frac{\delta}{2t}\right)^{3/2}.$$

By a union bound, we get that

$$\mathbb{P}\left(\exists t : |\xi_t| \geq \sqrt{3 \log(2t/\delta)}\right) \leq \sum_{t=1}^{\infty} \left(\frac{\delta}{2t}\right)^{3/2} \leq \delta^{3/2} < \delta.$$

□

Lemma 5. *For any (possibly randomized) player strategy, it holds that*

$$\mathbb{P}\left(\mathbb{E}_{\text{player}}[R_T] \geq \frac{1}{40}T^{2/3}\right) \geq \frac{1}{40},$$

where \mathbb{P} is over the adversary's randomization, and $\mathbb{E}_{\text{player}}[R_T]$ is the player's expected regret (over the player's randomization).

Proof. By Lemma 2, we already know that

$$\mathbb{E}[\mathbb{E}_{\text{player}}[R_T]] \geq \frac{1}{10}T^{2/3}, \quad (15)$$

since if we have a $T^{2/3}/10$ lower bound on the regret for any deterministic player strategy, the same holds for any randomized player strategy. Our approach is to apply Lemma 3 in order to convert this into a probability lower bound as in the lemma statement. However, we cannot apply Lemma 3 as-is, since $\mathbb{E}_{\text{player}}[R_T]$ can be as large as $\Omega(T)$, and the resulting bound is too weak. Instead, we show that there exists a *different* player strategy, with expected regret $\mathbb{E}_{\text{player}}^{\sim}[R_T]$, such that $|\mathbb{E}_{\text{player}}^{\sim}[R_T]|$ is always at most $2T^{2/3}$ and

$$\mathbb{E}_{\text{player}}^{\sim}[R_T] \leq 2 \mathbb{E}_{\text{player}}[R_T] \quad (16)$$

for any realization of the adversary's random strategy. Also, analogous to Eq. (15), we have $\mathbb{E}[\mathbb{E}_{\text{player}}^{\sim}[R_T]] \geq \frac{1}{10}T^{2/3}$ by Lemma 2. Therefore, using Eq. (16) and Lemma 3, we get that

$$\mathbb{P}\left(\mathbb{E}_{\text{player}}[R_T] \geq \frac{1}{40}T^{2/3}\right) \geq \mathbb{P}\left(\mathbb{E}_{\text{player}}^{\sim}[R_T] \geq \frac{1}{20}T^{2/3}\right) \geq \frac{1}{40}$$

as required.

The new player strategy we consider depends on the horizon T , and is very simple: It is identical to the original player strategy, but whenever the number of action switches reaches $\lfloor T^{2/3} \rfloor$, the player “freezes” in its current action, and keeps playing the same action till T rounds are elapsed. Clearly, the number of switches with this strategy can never be more than $T^{2/3}$, and since the regret in terms of the loss ℓ_t at each round is either 0 or $T^{-1/3}$, we get that the total regret R_T can never be more than $T^{2/3} + T * T^{-1/3} = 2T^{2/3}$.

To prove Eq. (16), we consider some instantiation of the adversary's random strategy, and note that for any realization of the player's random coin tosses, the regret can only differ between the two strategies if S_T (the total number of switches) is at least $\lfloor T^{2/3} \rfloor$. Therefore, we have $\mathbb{P}_{\text{player}}(S_T < \lfloor T^{2/3} \rfloor) = \mathbb{P}_{\text{player}}^{\sim}(S_T < \lfloor T^{2/3} \rfloor)$, $\mathbb{P}_{\text{player}}(S_T \geq \lfloor T^{2/3} \rfloor) = \mathbb{P}_{\text{player}}^{\sim}(S_T \geq \lfloor T^{2/3} \rfloor)$ and $\mathbb{E}_{\text{player}}[R_T | S_T < \lfloor T^{2/3} \rfloor] = \mathbb{E}_{\text{player}}^{\sim}[R_T | S_T < \lfloor T^{2/3} \rfloor]$. Also, we recall that $R_T \geq 0$ with the adversary strategy that we consider (since one action is always worse than the other action at all rounds). Finally, we note that if $S_T \geq \lfloor T^{2/3} \rfloor$, then the regret for both strategies is at least $\lfloor T^{2/3} \rfloor$ (since with the adversary strategy that we consider, the number of switches is a lower bound on the

regret). Using these observations, we have

$$\begin{aligned}
& \mathbb{E}_{\widetilde{\text{player}}}[R_T] \\
&= \mathbb{P}_{\widetilde{\text{player}}}(S_T < \lfloor T^{2/3} \rfloor) \mathbb{E}_{\widetilde{\text{player}}}[R_T | S_T < \lfloor T^{2/3} \rfloor] + \mathbb{P}_{\widetilde{\text{player}}}(S_T \geq \lfloor T^{2/3} \rfloor) \mathbb{E}_{\widetilde{\text{player}}}[R_T | S_T \geq \lfloor T^{2/3} \rfloor] \\
&\leq \mathbb{P}_{\widetilde{\text{player}}}(S_T < \lfloor T^{2/3} \rfloor) \mathbb{E}_{\widetilde{\text{player}}}[R_T | S_T < \lfloor T^{2/3} \rfloor] + \mathbb{P}_{\widetilde{\text{player}}}(S_T \geq \lfloor T^{2/3} \rfloor) 2T^{2/3} \\
&= \mathbb{P}_{\text{player}}(S_T < \lfloor T^{2/3} \rfloor) \mathbb{E}_{\text{player}}[R_T | S_T < \lfloor T^{2/3} \rfloor] + \mathbb{P}_{\text{player}}(S_T \geq \lfloor T^{2/3} \rfloor) 2T^{2/3} \\
&\leq 2 \left(\mathbb{P}_{\text{player}}(S_T < \lfloor T^{2/3} \rfloor) \mathbb{E}_{\text{player}}[R_T | S_T < \lfloor T^{2/3} \rfloor] + \mathbb{P}_{\text{player}}(S_T \geq \lfloor T^{2/3} \rfloor) T^{2/3} \right) \\
&\leq 2 \left(\mathbb{P}_{\text{player}}(S_T < \lfloor T^{2/3} \rfloor) \mathbb{E}_{\text{player}}[R_T | S_T < \lfloor T^{2/3} \rfloor] + \mathbb{P}_{\text{player}}(S_T \geq \lfloor T^{2/3} \rfloor) \mathbb{E}_{\text{player}}[R_T | S_T \geq \lfloor T^{2/3} \rfloor] \right) \\
&= 2 \mathbb{E}_{\text{player}}[R_T],
\end{aligned}$$

where in the second-to-last step we used the fact that if $S_T \geq \lfloor T^{2/3} \rfloor$, then the regret is at least $\lfloor T^{2/3} \rfloor$, plus we must have picked the worst action (worst by $T^{-1/3}$ than the best action) at least $\Omega(T^{2/3})$ times, hence the total regret is certainly at least $T^{2/3}$. \square

Finally, we use Lemma 4 with $\delta = 1/80$, to get that with probability at least $1 - 1/80$, the drift factor D_t of the adversarial strategy is at most $\sqrt{3 \log(160t)} \leq \sqrt{3 \log(t)} + 16$ for all t . Moreover, Lemma 5 tells us that $\mathbb{E}_{\text{player}}[R_T]$ is at least $\frac{1}{40} T^{2/3}$ with probability at least $1/40$. This implies that the intersection of the two events is non-empty, and there exists some deterministic adversarial strategy, such that the drift $D_t \leq \sqrt{3 \log(t)} + 16$ for all t , and the expected regret is at least $\frac{1}{40} T^{2/3}$ as required.

C Proof of Thm. 2

Thm. 1 guarantees that given any player's strategy, there is some deterministic adversary strategy with a lower bound on the regret. However, as part of proving Thm. 1, we actually showed that there exists some *randomized* adversary strategy $\{\hat{f}_t\}_{t=1}^T$ with memory size 1, such that for *any* (possibly randomized) player strategy $x_{1:t}$,

$$\mathbb{E} \left[\sum_{t=1}^T \hat{f}_t(X_{t-1}, X_t) - \min_{x \in \mathcal{A}} \sum_{t=1}^T \hat{f}_t(x, x) \right] \geq \frac{1}{10} T^{2/3} \quad (17)$$

(see Lemma 2). We now use this strategy to define a randomized adversary strategy for our setting (with memory size 2), for a game of $T + 1$ rounds. We let $f_1(x_1) = 0$ for any x_1 , $f_2(x_1, x_2) = \hat{f}_1(x_1)$, and for every $t = 3 \dots T + 1$,

$$f_t(x_{t-2}, x_{t-1}, x_t) = \hat{f}_{t-1}(x_{t-2}, x_{t-1}). \quad (18)$$

Now, suppose we had some (possibly randomized) player strategy $X_1 \dots X_{T+1}$, so that in expectation over the player and adversary strategies, we have

$$\mathbb{E} \left[\sum_{t=1}^{T+1} f_t(X_{t-2}, X_{t-1}, X_t) - \min_{x \in \mathcal{A}} \sum_{t=1}^{T+1} f_t(x, x, x) \right] < \frac{1}{10} T^{2/3}.$$

In particular, since f_1 is always 0, it would imply that

$$\mathbb{E} \left[\sum_{t=2}^{T+1} f_t(X_{t-2}, X_{t-1}, X_t) - \min_{x \in \mathcal{A}} \sum_{t=2}^{T+1} f_t(x, x, x) \right] < \frac{1}{10} T^{2/3}.$$

By Eq. (18), this implies

$$\mathbb{E} \left[\sum_{t=1}^T \hat{f}_t(X_{t-1}, X_t) - \min_{x \in \mathcal{A}} \sum_{t=1}^T \hat{f}_t(x, x) \right] < \frac{1}{10} T^{2/3}.$$

Thus, if we could implement the player strategy $X_1 \dots X_T$ in the bandits-with-switching-costs setting, it will contradict Eq. (17). To see that this indeed can happen, note that each X_t is a (possibly

randomized) function of $X_{1:t-1}$ as well as $\{f_\tau(X_{\tau-2}, X_{\tau-1}, X_\tau)\}_{\tau=1}^{t-1}$. But again, due to Eq. (18) and the fact that f_1 is always 0, X_t can in fact be defined using $X_{1:t-1}$ and

$$\{f_\tau(X_{\tau-2}, X_{\tau-1}, X_\tau)\}_{\tau=2}^{t-1} = \{\hat{f}_{\tau-1}(X_{\tau-2}, X_{\tau-1})\}_{\tau=2}^{t-1}.$$

The right hand side is an observable quantity in the bandit setting: In each round t , we know what are the set of losses $\{\hat{f}_{\tau-1}(X_{\tau-2}, X_{\tau-1})\}_{\tau=2}^{t-1}$ that we obtained. Thus, we can simulate the strategy $x_{1:t}$ in the bandit-with-switching-costs setting, and get an expected regret smaller than $\frac{1}{10}T^{2/3}$, contradicting Eq. (17). Thus, the expected regret (for a game of $T+1$ rounds) must be at least $\frac{1}{10}T^{2/3}$. Substituting T instead of $T+1$, we get that the expected regret for a game with T rounds is at least $\frac{1}{10}(T-1)^{2/3}$.

The regret bound we just now obtained is in expectation over the randomized adversary strategy, and holds for any player's strategy. We now use the same line of argument as in the last part of Thm. 1's proof, to show that for any (possibly randomized) player's strategy, there exists some *deterministic* adversary strategy, with a similar expected regret bound, and with losses of bounded drift. Specifically, a result completely analogous to Lemma 5 implies that

$$\mathbb{P}\left(\mathbb{E}_{\text{player}}[R_T] \geq \frac{1}{40}(T-1)^{2/3}\right) \geq \frac{1}{40}\left(\frac{T-1}{T}\right)^{2/3},$$

which is at least $1/80$ for any $T > 1$ (if $T = 1$ the bound in the theorem is trivial from the non-negativity of R_T for the adversary strategy that we consider). Moreover, using Eq. (4) as in the proof of Thm. 1, the probability of the loss drift being at most $\sqrt{3\log(320t)} \leq \sqrt{3\log(t) + 18}$ is at least $1 - 1/160$. Thus, the intersection of the two events is not empty, and this implies that there exists some *deterministic* adversary strategy causing expected regret $\geq \frac{1}{40}(T-1)^{2/3}$, and loss drift at most $\sqrt{3\log(t) + 18}$ for all t .

D Proofs of Upper Bounds

Proof of Thm. 3. Each loss functions equals $f_t(x_{1:t}) = \ell_t(x_t) + \mathbb{I}_{\{x_t \neq x_{t-1}\}}$, where ℓ_t is an oblivious loss function. Since the range of f_t is contained in an interval of size C , the range of ℓ_t must be contained in an interval of size $C-1$. In other words,

$$\forall x \in \mathcal{A} \quad \ell_t(x) - \min_y \ell_t(y) \leq C-1.$$

Therefore, by definition, the range of ℓ'_t is contained in the interval $[0, 1]$, and the analysis of the FLL algorithm holds. Namely, if X_1, X_2, \dots is the sequence of actions chosen by FLL, then, for any T

$$\mathbb{E}\left[\sum_{t=1}^T \ell'_t(X_t)\right] - \min_{x \in \mathcal{A}} \sum_{t=1}^T \ell'_t(x) = \mathcal{O}(\sqrt{T}), \quad (19)$$

and

$$\mathbb{E}\left[\sum_{t=1}^T \mathbb{I}_{\{X_t \neq X_{t-1}\}}\right] = \mathcal{O}(\sqrt{T}). \quad (20)$$

Plugging the definition of ℓ'_t into Eq. (19) and rearranging terms, we get

$$\mathbb{E}\left[\sum_{t=1}^T \ell_t(X_t)\right] - \min_{x \in \mathcal{A}} \sum_{t=1}^T \ell_t(x) = (C-1)\mathcal{O}(\sqrt{T}).$$

Summing the above with Eq. (20) gives

$$\mathbb{E}\left[\sum_{t=1}^T f_t(X_{1:t})\right] - \min_{x \in \mathcal{A}} \sum_{t=1}^T f_t(x \dots x) = \mathcal{O}(C\sqrt{T}).$$

□

Proof of Thm. 4. Recall that $f'_t(x_{1:t}) = \frac{1}{2(C+D)} (\ell_t(x_t) - \ell_{t-1}(x_{t-1})) + \frac{1}{2}$, and note that our assumptions imply that

$$\begin{aligned} |\ell_t(x_t) - \ell_{t-1}(x_{t-1})| &= |\ell_t(x_t) - \ell_{t-1}(x_t) + \ell_{t-1}(x_t) - \ell_{t-1}(x_{t-1})| \\ &\leq |\ell_t(x_t) - \ell_{t-1}(x_t)| + |\ell_{t-1}(x_t) - \ell_{t-1}(x_{t-1})| \\ &\leq D + C . \end{aligned}$$

Therefore, $f'_t(x_{1:t})$ is always bounded in $[0, 1]$. Moreover, the action x which minimizes $\sum_{t=1}^T f'_t(X_{1:t-1}, x)$ does not depend on the player's actions (only on $\sum_{t=1}^T \ell_t(x)$). Letting x^* denote such an action, the standard analysis for Exp3 implies that

$$\mathbb{E} \left[\sum_{t=1}^T f'_t(X_{1:t}) - \min_{x \in \mathcal{A}} \sum_{t=1}^T f'_t(X_{1:t-1}, x) \right] = \mathbb{E} \left[\sum_{t=1}^T f'_t(X_{1:t}) - \sum_{t=1}^T f'_t(X_{1:t-1}, x^*) \right] = \mathcal{O}(\sqrt{T}),$$

where $X_{1:T}$ is the sequence of actions chosen by Exp3. Using the definition of f'_t , the left hand side above can be rewritten as

$$\begin{aligned} &\frac{1}{2(C+D)} \mathbb{E} \left[\sum_{t=1}^T (\ell_t(X_t) - \ell_{t-1}(X_{t-1})) - \min_{x \in \mathcal{A}} \sum_{t=1}^T (\ell_t(x) - \ell_{t-1}(X_{t-1})) \right] \\ &= \frac{1}{2(C+D)} \mathbb{E} \left[\sum_{t=1}^T \ell_t(X_t) \right] - \min_{x \in \mathcal{A}} \sum_{t=1}^T \ell_t(x) . \end{aligned}$$

Therefore,

$$\mathbb{E}[R_T] = \mathbb{E} \left[\sum_{t=1}^T \ell_t(X_t) \right] - \min_{x \in \mathcal{A}} \sum_{t=1}^T \ell_t(x) = 2(C+D)\mathcal{O}(\sqrt{T}) .$$

Using the assumption that $D_t = \mathcal{O}(\sqrt{\log(T)})$, we conclude that $\mathbb{E}[R_T] = \tilde{\mathcal{O}}(C\sqrt{T})$. \square

Proof of Thm. 5. First, note that, due to the bounded range and drift assumptions, $\hat{f}_t \in [0, 1]$. Also note that

$$f_t(x_{t-m:t}) - f_t(x \dots x) = 2(C + (m+1)D)(\hat{f}_t(x_{t-m:t}) - \hat{f}_t(x \dots x)) .$$

As previously mentioned, we divide the T rounds into J consecutive epochs of the same length T/J , where $T/J \geq 2K(m+1)$, plus an additional final epoch of length at most T/J . We let t_j denote the index of the first round in the j -th epoch. We run a mini-batched version of the Hedge algorithm [11] over the epochs: at the beginning of each epoch j , Hedge draws an action $X_j \in \mathcal{A}$ which is played consistently throughout the epoch. Now assume that at the end of each epoch j , loss estimates $g_j(x) \in [0, 1]$ for each action x are available such that

$$\mathbb{E}[g_j(x)] = \frac{1}{T/J - 2m - 1} \sum_{t=t_j+2m+1}^{t_{j+1}-1} \hat{f}_t(x \dots x)$$

where the randomness used to compute each g_j is independent of that used by Hedge to draw X_j . At the end of epoch j , we feed loss estimates $g_j(x)$ for each $x \in \mathcal{A}$ to Hedge. The resulting regret

can be bounded as follows,

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} [f_t(X_{t-m:t}) - f_t(x \dots x)] \\
& \leq \sum_{j=1}^J \sum_{t=t_j}^{t_{j+1}-1} \mathbb{E} [f_t(X_{t-m:t}) - f_t(x \dots x)] + \frac{CT}{J} \\
& = 2(C + (m+1)D) \sum_{j=1}^J \sum_{t=t_j}^{t_j+2m} \mathbb{E} [\hat{f}_t(X_{t-m:t}) - \hat{f}_t(x \dots x)] \\
& \quad + 2(C + (m+1)D) \sum_{j=1}^J \sum_{t=t_j+2m+1}^{t_{j+1}-1} \mathbb{E} [\hat{f}_t(X_j \dots X_j) - \hat{f}_t(x \dots x)] + \frac{CT}{J} \\
& \leq 2(C + (m+1)D)(2m+1)J \\
& \quad + 2(C + (m+1)D) \frac{T}{J} \mathbb{E} \left[\sum_{j=1}^J \mathbb{E} [g_j(X_j) - g_j(x) | X_j] \right] + \frac{CT}{J} \\
& = 2(C + (m+1)D)(2m+1)J \\
& \quad + 2(C + (m+1)D) \frac{T}{J} \mathbb{E} \left[\sum_{j=1}^J (g_j(X_j) - g_j(x)) \right] + \frac{CT}{J} \\
& \leq 2(C + (m+1)D)(2m+1)J + 4(C + (m+1)D) \frac{T}{J} \sqrt{J \ln K} + \frac{CT}{J}.
\end{aligned}$$

In the last step we applied the known upper bound on the regret of Hedge with respect to losses $g_j \in [0, 1]$, where K is the number of actions. This is valid if, in particular, losses g_j are oblivious. We now explain how to obtain oblivious estimates g_j with the desired properties. At the beginning of each epoch j , we use the independent randomization to draw K exploration steps $\{t_x : x \in \mathcal{A}\}$ from the set $T_j = \{t_j, \dots, t_{j+1} - 2m - 2\}$ with the property that these steps are well separated. Namely, between any two t_x and $t_{x'}$ there are at least $2m + 1$ consecutive free time steps in T_j . During epoch j , when we arrive at step t_x we freeze Hedge and play action x_0 for $m + 1$ time steps, then we play action x for $m + 1$ more time steps. We use the two observed losses $f_{t_x+m}(x_0 \dots x_0)$ and $f_{t_x+2m+1}(x \dots x)$ to compute $\hat{f}_{t_x+2m+1}(x \dots x)$. Because the t_x are well separated, the exploration steps do not interfere with each other. Suppose now that we can draw these points such that the marginal of each t_x is uniform in T_j . Then

$$\begin{aligned}
\mathbb{E} [\hat{f}_{t_x+2m+1}(x \dots x)] &= \frac{1}{T/J - 2m - 1} \sum_{t=t_j}^{t_{j+1}-2m-2} \hat{f}_{t+2m+1}(x \dots x) \\
&= \frac{1}{T/J - 2m - 1} \sum_{t=t_j+2m+1}^{t_{j+1}-1} \hat{f}_t(x \dots x).
\end{aligned}$$

This shows that $\hat{f}_{t_x+2m+1}(x \dots x)$ is a valid estimate $g_j(x)$. Moreover, for each $x \in \mathcal{A}$ the quantity $\hat{f}_{t_x+2m+1}(x \dots x)$ does not depend on Hedge's action X_j for the current epoch j . It does not even depend on Hedge's past actions. Hence, Hedge is indeed run on a set of oblivious losses and the standard regret bound applies.

The last thing to prove is that we can draw $\{t_x : x \in \mathcal{A}\} \subset T_j$ such that the marginal of each t_x is uniform in T_j . Note that giving equal probability to all well separated configurations of $\{t_x : x \in \mathcal{A}\}$ does not work, because the times steps closer to the beginning and to the end of T_j appear in more configurations (for example, check the case $|T_j| = 8$ and $m = 1$). This problem can be fixed simply by arranging the points of T_j on a circle, so that the first point t_j follows the last point $t_{j+1} - 2m - 2$, and then enforcing well-separatedness on the circle. This makes the sample space completely symmetric, excluding those configurations of exploration points that exploited border effects.

The additional regret due to the computation of the K exploration points is $2(m+1)CK$ per epoch. The final regret, including these additional costs, is then bounded by

$$2(C + (m+1)D)(2m+1)J + 4(C + (m+1)D)\frac{T}{J}\sqrt{J\ln K} + \frac{CT}{J} + 2(m+1)CKJ.$$

Choosing J of order $T^{2/3}$ concludes the proof. \square