# On First-Order Bounds, Variance and Gap-Dependent Bounds for Adversarial Bandits 

Roman Pogodin<br>Gatsby Computational Neuroscience Unit<br>University College London, London, UK<br>roman.pogodin.17@ucl.ac.uk

Tor Lattimore<br>DeepMind<br>London, UK<br>tor.lattimore@gmail.com


#### Abstract

We make three contributions to the theory of $k$-armed adversarial bandits. First, we prove a first-order bound for a modified variant of the INF strategy by Audibert and Bubeck [2009], without sacrificing worst case optimality or modifying the loss estimators. Second, we provide a variance analysis for algorithms based on follow the regularised leader, showing that without adaptation the variance of the regret is typically $\Omega\left(n^{2}\right)$ where $n$ is the horizon. Finally, we study bounds that depend on the degree of separation of the arms, generalising the results by Cowan and Katehakis [2015] from the stochastic setting to the adversarial and improving the result of Seldin and Slivkins [2014] by a factor of $\log (n) / \log \log (n)$.


## 1 INTRODUCTION

The $k$-armed adversarial bandit is a sequential game played over $n$ rounds. At the start of the game the adversary secretly chooses a sequence of losses $\left(\ell_{t}\right)_{t=1}^{n}$ with $\ell_{t} \in[0,1]^{k}$. In each round $t$ the learner chooses a distribution $P_{t}$ over the actions $[k]=\{1,2, \ldots, k\}$. An action $A_{t} \in[k]$ is sampled from $P_{t}$ and the learner observes the loss $\ell_{t A_{t}}$. Like prior work we focus on controlling the regret, which is

$$
\hat{R}_{n}=\max _{i \in[k]} \sum_{t=1}^{n}\left(\ell_{t A_{t}}-\ell_{t i}\right) .
$$

This quantity is a random variable, so the standard objective is to bound $\hat{R}_{n}$ with high probability or its expectation: $R_{n}=\mathbb{E}\left[\hat{R}_{n}\right]$.
We make three contributions, with the common objective of furthering our understanding of the application of
follow the regularised leader (FTRL) to adversarial bandit problems. Our first contribution is a modification of the INF policy by Audibert and Bubeck [2009] in order to prove first-order bounds (i.e. in terms of the loss of the best action) without sacrificing minimax optimality. Then we turn our attention to the variance of algorithms based on FTRL. Here we prove that using the standard importance-weighted estimators and a large class of potentials leads to a variance of $\Omega\left(n^{2}\right)$, which is the worst possible for bounded losses. Finally, we investigate the asymptotic performance of algorithms when there is a linear separation between the losses of the arms. We improve the result by Seldin and Slivkins [2014] by a factor of $\log (n) / \log \log (n)$ and generalise known results in the stochastic setting by Cowan and Katehakis [2015] to the adversarial one by constructing an algorithm for which the regret grows arbitrarily slowly almost surely.

Related work The literature on adversarial bandits is enormous. See the books by Bubeck and Cesa-Bianchi [2012] and Lattimore and Szepesvári [2019] for a comprehensive account. The common thread in the three components of our analysis is adaptivity for algorithms based on follow the regularised leader. The INF policy that underlies much of our analysis was introduced by Audibert and Bubeck [2009]. The connection to mirror descent and follow the regularised leader came later [Audibert and Bubeck, 2010, Bubeck and CesaBianchi, 2012], which greatly simplified the analysis. The principle justification for introducing this algorithm was to prove bounds on the minimax regret. Remarkably, it was recently shown that by introducing a nonadaptive decaying learning rate, the algorithm retains minimax optimality while simultaneously achieving a near-optimal logarithmic regret in the stochastic setting [Zimmert and Seldin, 2019]. Despite its simplicity, the algorithm improves on the state-of-the-art for this problem Bubeck and Slivkins [2012], Seldin and Slivkins [2014], Seldin and Lugosi [2017]. See also
the extension to the combinatorial semibandit setting [Zimmert et al., 2019]. First-order bounds for bandits were first given by Allenberg et al. [2006], who analysed a modification of Exp3 [Auer et al., 1995]. As far as we know, previous algorithms with first order bounds have not been minimax optimal $\left(R_{n}=O(\sqrt{k n})\right.$ ): the recent work by Neu [2015b] achieved $O(\sqrt{k n(\log (k)+1)})$ expected regret, and [Wei and Luo, 2018] had a $O(\sqrt{k n \log n})$ bound. Both papers used the idea of an adaptive learning rate similar to our analysis. In the setting of gains rather than losses Audibert and Bubeck [2010] have shown that by introducing biased estimators it is possible to prove a bound of $O\left(\sqrt{k G^{*}}\right)$ where $G^{*}$ is the maximum gain. Although it is not obvious, we suspect the same idea could be applied in our setting. We find it interesting nevertheless that the same affect is possible without modifying the loss estimators. The aforementioned work also assumes knowledge of $G^{*}$. Possibly our adaptive learning rates could be used to make this algorithm anytime without a doubling trick.

Although it is well known that straightforward applications of follow the regularised leader or mirror descent with importance-weighted estimators leads to poor concentration of the regret, we suspect the severity of the situation is not widely appreciated. As far as we know, the quadratic variance of Exp3 was only derived recently [Lattimore and Szepesvári, 2019, §11]. There are, however, a number of works modifying the importance-weighted estimators to prove high probability bounds Auer et al. [1995], Abernethy and Rakhlin [2009], Neu [2015a] with matching lower bounds by Gerchinovitz and Lattimore [2016]. Finally, we note there are many kinds of adaptivity beyond first-order bounds. For example sparsity and variance [Bubeck et al., 2018, Hazan and Kale, 2011, and others].

## 2 NOTATION

Given a vector $x \in \mathbb{R}^{d}$ let $\operatorname{diag}(x) \in \mathbb{R}^{d \times d}$ be the diagonal matrix with $x$ along the diagonal. The interior of a topological space $X$ is interior $(X)$ and its boundary is $\partial X$. The standard basis vectors are $e_{1}, \ldots, e_{d}$. The ( $d-1$ )-dimensional probability simplex is $\Delta^{d-1}=\{x \in$ $\left.[0,1]^{d}:\|x\|_{1}=1\right\}$. A convex function $F: \mathbb{R}^{d} \rightarrow \mathbb{R} \cup$ $\{\infty\}$ has domain $\operatorname{dom}(F)=\left\{x \in \mathbb{R}^{d}: F(x) \neq \infty\right\}$. The Bregman divergence with respect to a differentiable $F$ is a function $D_{F}: \operatorname{dom}(F) \times \operatorname{dom}(F) \rightarrow[0, \infty]$ defined by $D_{F}(x, y)=F(x)-F(y)-\langle\nabla F(y), x-y\rangle$. The Fenchel dual of $F$ is $F^{*}: \mathbb{R}^{d} \rightarrow \mathbb{R} \cup\{\infty\}$ defined by $F^{*}(u)=\sup _{x \in \mathbb{R}^{d}}\langle x, u\rangle-F(x)$.
There are $k$ arms and the horizon is $n$, which may or may not be known. The losses are $\left(\ell_{t}\right)_{t=1}^{n}$ with $\ell_{t} \in$ $[0,1]^{k}$. We let $L_{t}=\sum_{s=1}^{t} \ell_{s}$. The importance-weighted

| Potential | Definition | Alg. |
| :--- | :--- | :--- |
| Negentropy | $\frac{1}{\eta} \sum_{i=1}^{k} p_{i}\left(\log \left(p_{i}\right)-1\right)$ | Exp3 |
| $1 / 2$-Tsallis | $-\frac{2}{\eta} \sum_{i=1}^{k} \sqrt{p_{i}}$ | INF |
| Log barrier | $-\frac{1}{\eta} \sum_{i=1}^{k} \log \left(p_{i}\right)$ |  |

Table 1: Common potential functions
estimator of $\ell_{t}$ is $\hat{\ell}_{t}$ defined by $\hat{\ell}_{t i}=\mathbb{1}\left\{A_{t}=i\right\} \ell_{t i} / P_{t i}$. All algorithms proposed here ensure that $P_{t i}>0$ for all $t$ and $i$, so this quantity is always well defined. Let $\hat{L}_{t}=\sum_{s=1}^{t} \hat{\ell}_{s}$. Expectations are with respect to the randomness in the actions $\left(A_{t}\right)_{t=1}^{n}$. Of course the learner can only choose $P_{t}$ based on information available at the start of round $t$. Let $\mathcal{F}_{t}=\sigma\left(A_{1}, \ldots, A_{t}\right)$. Then $P_{t}$ is $\mathcal{F}_{t-1}$-measurable. Let $A_{t i}=\mathbb{1}\left\{A_{t}=i\right\}$ and $T_{i}(t)=\sum_{s=1}^{t} A_{s i}$ be the number of times arm $i$ is played in the first $t$ rounds. Our standing assumption is that the first arm is optimal. All our algorithms are symmetric, so this is purely for notational convenience.
Assumption 2.1. $L_{t 1}=\min _{i \in[k]} L_{t i}$.

## 3 FOLLOW THE REGULARISED LEADER

Follow the regularized leader (FTRL) is a popular tool for online optimization [Shalev-Shwartz, 2007, Hazan, 2016]. The basic algorithm depends on a sequence of potential functions $\left(F_{t}\right)_{t=1}^{\infty}$ where $F_{t}: \mathbb{R}^{d} \rightarrow \mathbb{R} \cup\{\infty\}$ is convex and $\operatorname{dom}\left(F_{t}\right) \cap \Delta^{k-1} \neq \emptyset$. In each round the algorithm chooses the distribution

$$
P_{t}=\underset{p \in \Delta^{k-1}}{\arg \min }\left\langle p, \hat{L}_{t-1}\right\rangle+F_{t}(p)
$$

which we assume exists. The action $A_{t} \in[k]$ is sampled from $P_{t}$. In many applications $F_{t}=F$ is chosen in a time independent way, with examples given in Table 1. This has the disadvantage that $F$ must be chosen in advance in a way that depends on the horizon, which may be unknown. This weakness can be overcome by choosing $F_{t}=F / \eta_{t}$ where $\left(\eta_{t}\right)_{t=1}^{\infty}$ is a sequence of learning rates, which may be chosen in advance or adaptively in a data-dependent way.
A modification that will prove useful is to let $\left(\mathcal{A}_{t}\right)_{t=1}^{\infty}$ be a sequence of subsets of $\Delta^{k-1}$ and define

$$
P_{t}=\underset{p \in \mathcal{A}_{t}}{\arg \min }\left\langle p, \hat{L}_{t-1}\right\rangle+F_{t}(p)
$$

The restriction to a subset of $\Delta^{k-1}$ can be useful to control the gradients of $F_{t}\left(P_{t}\right)$, which is sometimes crucial. The following theorem provides a generic bound for FTRL with changing potentials and constraint sets.

The result is reminiscent of many previous bounds for FTRL, but a reference for this result seems elusive. Most related is the generic analysis by Joulani et al. [2017], which also provides the most comprehensive literature summary.

Theorem 3.1. Assume $\mathcal{A}_{1} \subseteq \cdots \subseteq \mathcal{A}_{n+1} \subseteq \Delta^{k-1}$ and $\left(F_{t}\right)_{t=1}^{n+1}$ is a sequence of convex functions with $\operatorname{dom}\left(F_{t}\right) \cap \mathcal{A}_{t} \neq \emptyset$ for all $t$. Define

$$
\begin{gathered}
d_{t}=\max _{y \in \mathcal{A}_{t+1}} \min _{x \in \mathcal{A}_{t}}\|x-y\|_{1}, \quad g_{t}=\sup _{x \in \mathcal{A}_{t}}\left\|\nabla F_{t}(x)\right\|_{\infty} \\
\text { and } \quad v_{n}=\sum_{t=1}^{n} d_{t}\left(g_{t}+(t-1)\right) .
\end{gathered}
$$

Then the regret of FTRL is bounded by

$$
\begin{aligned}
R_{n} \leq & v_{n}+\mathbb{E}\left[\sum_{t=1}^{n}\left\langle P_{t}-P_{t+1}, \hat{\ell}_{t}\right\rangle-D_{F_{t}}\left(P_{t+1}, P_{t}\right)\right] \\
& +\mathbb{E}\left[\min _{p \in \mathcal{A}_{n+1}}\left(F_{n+1}(p)+n\left\|p-e_{1}\right\|_{1}\right)-F_{1}\left(P_{1}\right)\right] \\
& +\mathbb{E}\left[\sum_{t=1}^{n}\left(F_{t}\left(P_{t+1}\right)-F_{t+1}\left(P_{t+1}\right)\right)\right] .
\end{aligned}
$$

Proof. Let $p \in \mathcal{A}_{n+1}$. Using the fact that $\hat{\ell}_{t}$ is unbiased,

$$
\begin{aligned}
R_{n} & =\mathbb{E}\left[\sum_{t=1}^{n}\left\langle P_{t}-e_{1}, \hat{\ell}_{t}\right\rangle\right] \\
& =\mathbb{E}\left[\sum_{t=1}^{n}\left\langle P_{t}-p, \hat{\ell}_{t}\right\rangle\right]+\mathbb{E}\left[\sum_{t=1}^{n}\left\langle p-e_{1}, \ell_{t}\right\rangle\right] .
\end{aligned}
$$

The second sum is the approximation error, and by Holder's inequality,

$$
\sum_{t=1}^{n}\left\langle p-e_{1}, \ell_{t}\right\rangle \leq\left\|p-e_{1}\right\|_{1} \sum_{t=1}^{n}\left\|\ell_{t}\right\|_{\infty} \leq n\left\|p-e_{1}\right\|_{1}
$$

Therefore,

$$
\begin{aligned}
R_{n} \leq \mathbb{E}[ & \left.\sum_{t=1}^{n}\left\langle P_{t}-P_{t+1}, \hat{\ell}_{t}\right\rangle+\sum_{t=1}^{n}\left\langle P_{t+1}-p, \hat{\ell}_{t}\right\rangle\right] \\
& +n\left\|p-e_{1}\right\|_{1} .
\end{aligned}
$$

Let $\Phi_{t}(q)=F_{t}(q)+\sum_{s=1}^{t-1}\left\langle q, \hat{\ell}_{s}\right\rangle$, which is chosen so that $P_{t}=\arg \min _{q \in \mathcal{A}_{t}} \Phi_{t}(q)$. Then the second sum in
the above display equals

$$
\begin{aligned}
& \sum_{t=1}^{n}\left(\Phi_{t+1}\left(P_{t+1}\right)-\Phi_{t}\left(P_{t+1}\right)-F_{t+1}\left(P_{t+1}\right)+F_{t}\left(P_{t+1}\right)\right) \\
& \quad-\Phi_{n+1}(p)+F_{n+1}(p) \\
& =\sum_{t=1}^{n}\left(\Phi_{t}\left(P_{t}\right)-\Phi_{t}\left(P_{t+1}\right)\right) \\
& \quad+\Phi_{n+1}\left(P_{n+1}\right)-\Phi_{1}\left(P_{1}\right)-\Phi_{n+1}(p) \\
& \quad+F_{n+1}(p)+\sum_{t=1}^{n}\left(F_{t}\left(P_{t+1}\right)-F_{t+1}\left(P_{t+1}\right)\right)
\end{aligned}
$$

We can rewrite the $\Phi$-differences as

$$
\begin{aligned}
& \Phi_{t}\left(P_{t}\right)-\Phi_{t}\left(P_{t+1}\right) \\
& \quad=-D_{\Phi_{t}}\left(P_{t+1}, P_{t}\right)-\left\langle\nabla \Phi_{t}\left(P_{t}\right), P_{t+1}-P_{t}\right\rangle
\end{aligned}
$$

Let $\delta_{t}=P_{t+1}-\arg \min _{q \in \mathcal{A}_{t}}\left\|q-P_{t+1}\right\|_{1}$. Then due to first-order optimality condition for $P_{t}$ on $\mathcal{A}_{t}$,

$$
\mathbb{E}\left[\left\langle\nabla \Phi_{t}\left(P_{t}\right),\left(P_{t+1}-\delta_{t}\right)-P_{t}\right\rangle\right] \geq 0
$$

therefore

$$
\begin{aligned}
& \mathbb{E}\left[\left\langle\nabla \Phi_{t}\left(P_{t}\right), P_{t+1}-P_{t}\right\rangle\right] \geq \mathbb{E}\left[\left\langle\nabla \Phi_{t}\left(P_{t}\right), \delta_{t}\right\rangle\right] \\
& \quad \geq \mathbb{E}\left[\left\langle\nabla F_{t}\left(P_{t}\right), \delta_{t}\right\rangle+\sum_{s=1}^{t-1}\left\langle\hat{\ell}_{s}, \delta_{t}\right\rangle\right] \\
& \quad \geq-\mathbb{E}\left[\left\|\delta_{t}\right\|_{1}\left(\left\|\nabla F_{t}\left(P_{t}\right)\right\|_{\infty}+\sum_{s=1}^{t-1}\left\|\ell_{s}\right\|_{\infty}\right)\right] \\
& \quad \geq-d_{t} g_{t}-d_{t} \sum_{s=1}^{t-1}\left\|\ell_{s}\right\|_{\infty} \geq-d_{t}\left(g_{t}+(t-1)\right)
\end{aligned}
$$

where we used Holder's inequality, the definitions of $d_{t}$ and $g_{t}$, non-negativity of $\hat{\ell}_{s}$ and that $\mathbb{E} \hat{\ell}_{s}=\ell_{s} \in[0,1]$. It follows that
$\Phi_{t}\left(P_{t}\right)-\Phi_{t}\left(P_{t+1}\right) \leq d_{t}\left(g_{t}+k(t-1)\right)-D_{\Phi_{t}}\left(P_{t+1}, P_{t}\right)$.
Since $p \in \mathcal{A}_{n+1}$ and $P_{n+1}$ is the minimiser of $\Phi_{n+1}$ in $\mathcal{A}_{n+1}$, we have

$$
\Phi_{n+1}\left(P_{n+1}\right)-\Phi_{n+1}(p) \leq 0
$$

Finally, noting that $\Phi_{1}=F_{1}$ and $D_{\Phi_{t}}\left(P_{t+1}, P_{t}\right)=$ $D_{F_{t}}\left(P_{t+1}, P_{t}\right)$ we obtain

$$
\begin{aligned}
R_{n} \leq & n\left\|p-e_{1}\right\|_{1}+\sum_{t=1}^{n} d_{t}\left(g_{t}+(t-1)\right) \\
& +\mathbb{E}\left[\sum_{t=1}^{n}\left\langle P_{t}-P_{t+1}, \hat{\ell}_{t}\right\rangle-D_{F_{t}}\left(P_{t+1}, P_{t}\right)\right] \\
& +\mathbb{E}\left[F_{n+1}(p)-F_{1}\left(P_{1}\right)\right] \\
& +\mathbb{E}\left[\sum_{t=1}^{n}\left(F_{t}\left(P_{t+1}\right)-F_{t+1}\left(P_{t+1}\right)\right)\right]
\end{aligned}
$$

from which the statement follows.

## 4 FIRST ORDER BOUNDS

We now introduce the modification of the INF strategy, which takes inspiration from Wei and Luo [2018], Zimmert and Seldin [2019], Zimmert et al. [2019]. The new algorithm plays on the 'chopped' simplex, with the magnitude of the cut dependent on the round,

$$
\begin{equation*}
\mathcal{A}_{t}=\Delta^{k-1} \cap[1 / t, 1]^{k} \tag{1}
\end{equation*}
$$

Then for a convex potential $f_{t}(p)$ with $\operatorname{dom}\left(f_{t}\right)^{k} \cap$ $\Delta^{k-1} \neq \emptyset$ define a potential

$$
\begin{equation*}
F_{t}(p)=\frac{1}{\eta_{t}} \sum_{i=1}^{k} f_{t}\left(p_{i}\right) \tag{2}
\end{equation*}
$$

where the learning rate $\eta_{t}$ is given by

$$
\begin{equation*}
\eta_{t}=\frac{\eta_{0}}{\sqrt{1+\sum_{s=1}^{t-1} \hat{\ell}_{s A_{s}}^{2}\left(\nabla^{2}\left(f_{s}\right)\left(P_{s A_{s}}\right)\right)^{-1}}} \tag{3}
\end{equation*}
$$

where $\eta_{0}$ is positive constant to be tuned later.
The Hessian of the potential plays a fundamental role in the regret, simplifying the derivation of a generic firstorder bound:
Theorem 4.1. Suppose that $\nabla^{2} f_{t}$ is decreasing on $(0,1)$ and there exist $B, C \geq 0$ such that

$$
\frac{1}{p^{2} \nabla^{2} f_{t}(p)} \leq B, \quad \mathbb{E}\left[\frac{1}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}\right] \leq C
$$

for all $p \in(0,1)$ and $t \in[n]$. Assume additionally that there exist a non-negative constant $h_{1}$ and a nonnegative function $h_{2}(n)$ such that

$$
\begin{aligned}
v_{n} & +\min _{p \in \mathcal{A}_{n+1}}\left(F_{n+1}(p)+\left\|p-e_{1}\right\|_{1} n\right)-F_{1}\left(P_{1}\right) \\
& +\sum_{t=1}^{n}\left(F_{t}\left(P_{t+1}\right)-F_{t+1}\left(P_{t+1}\right)\right) \leq \frac{h_{1}}{\eta_{n+1}}+h_{2}(n),
\end{aligned}
$$

almost surely. Then the expected regret of FTRL with $\eta_{0}=\sqrt{h_{1}} / 2^{1 / 4}$ simultaneously satisfies

$$
\begin{aligned}
R_{n} & \leq \frac{\sqrt{h_{1}}}{2^{5 / 4}} B+h_{2}(n)+2 \sqrt{2} B h_{1}+2^{7 / 4} \sqrt{h_{1}} \\
& \times \sqrt{1+\frac{B L_{n 1}}{2}+\frac{B h_{2}(n)}{2}+B^{2}\left(\frac{\sqrt{h_{1}}}{2^{9 / 4}}+\frac{h_{1}}{\sqrt{2}}\right)} \\
R_{n} & \leq \frac{\sqrt{h_{1}}}{2^{5 / 4}} B+h_{2}(n)+2^{7 / 4} \sqrt{h_{1}} \sqrt{1+\frac{C n}{2}} .
\end{aligned}
$$

Remark 4.2. $h_{1}$ and $h_{2}(n)$ reflect the approximation error, non-stationarity of the potential $f_{t}$ and how sensitive it is to the changes in $\mathcal{A}_{t}$. In a simple case with $\mathcal{A}_{t}=\mathcal{A}, f_{t}=f$ for all $t$, this is a standard bound for the sum of the potential differences. $h_{1}$ can be a function of $n$ when the horizon $n$ is known, as we choose the learning rate based on it.

As an application of this general first-order result, we derive a worst-case optimal bound for a carefully chosen mixture of the INF regularizer and the log-barrier:
Corollary 4.3. For $\eta_{0}=k^{1 / 4} \sqrt{\frac{13}{3 \sqrt{2}}+\frac{3}{\sqrt{2} q}}$ and

$$
f_{t}(p)=-2 \sqrt{p}-\frac{\log p}{\sqrt{k} \log ^{1+q} \max \{3, t\}}
$$

and any $q>0$ and $n \geq 3$, the regret grows with $n$ as
$R_{n}=O\left(\sqrt{k L_{n 1} \log ^{1+q}(n)+k^{2} \log ^{2(1+q)} n}+k \log ^{1+q}(n)\right)$,
with some constants proportional to $1 / q$.
Corollary 4.4. For $q=1, \eta_{0}=k^{1 / 4} \sqrt{22 /(3 \sqrt{2})}$,

$$
\begin{aligned}
& R_{n} \leq 19 k^{2}+22 k \log ^{2}(n)+2 k \log (n)+6.5 \log (n) \\
& \quad \times \sqrt{k L_{n 1}+19 k^{3}+2 k^{2} \log (n)+11.2 k^{2} \log ^{2}(n)} .
\end{aligned}
$$

In the worst-case scenario the regret satisfies

$$
\limsup _{n \rightarrow \infty} \frac{R_{n}}{\sqrt{k n}} \leq 9.2
$$

Corollary 4.5. If the horizon $n \geq 3$ is known in advance, using $\mathcal{A}_{t}=\Delta^{k-1} \cap[1 / n, 1]^{k}, \eta_{0}=k^{1 / 4} \sqrt{3} / 2^{1 / 4}$ and

$$
f_{t}(p)=-2 \sqrt{p}-\frac{\log p}{\sqrt{k} \log n}
$$

results in

$$
\begin{aligned}
& R_{n} \leq \\
& \quad+4+9.1 k \log n \\
& \quad+4 \sqrt{k L_{n 1} \log (n)+2 \sqrt{k}+6 k^{2} \log ^{2}(n)} \\
& \quad \limsup _{n \rightarrow \infty} \frac{R_{n}}{\sqrt{k n}} \leq 5.9
\end{aligned}
$$

The proof of the last corollary simply repeats previous statements, also using the stationarity of the constraint set and $f_{t}(p)$. See the supplementary material for more details.

Remark 4.6. Theorem 4.1 with known $n$ reproduces the result of [Wei and Luo, 2018] (note that they used a slightly different algorithm and the learning rate schedule): for the $\log$-barrier potential $f_{t}(p)=-\log p$ we have $B=C=1$ and $h_{1}(n) \propto k \log n$, such that the worst-case regret is $R_{n}=O(\sqrt{k n \log n})$.

The proof of Theorem 4.1 follows from Theorem 3.1 and the following lemmas:

Lemma 4.7. For a potential of the form Eq. (2) with $\nabla^{2} f_{t}(p)$ that is monotonically decreasing on $p \in(0,1)$,

$$
\begin{aligned}
\sum_{t=1}^{n}\left\langle P_{t}-\right. & \left.P_{t+1}, \hat{\ell}_{t}\right\rangle-D_{F_{t}}\left(P_{t+1}, P_{t}\right) \\
& \leq \sum_{t=1}^{n} \frac{\eta_{t}}{2} \frac{\ell_{t A_{t}}^{2}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}
\end{aligned}
$$

Proof. Let $t \in[n]$ and suppose that $P_{t+1, A_{t}}>P_{t A_{t}}$. Then using the fact that the loss estimators and Bregman divergence are non-negative,

$$
\begin{aligned}
& \left\langle P_{t}-P_{t+1}, \hat{\ell}_{t}\right\rangle-D_{F_{t}}\left(P_{t+1}, P_{t}\right) \leq\left\langle P_{t}-P_{t+1}, \hat{\ell}_{t}\right\rangle \\
& =\left(P_{t A_{t}}-P_{\left.t+1, A_{t}\right)}\right) \hat{\ell}_{t A_{t}} \leq 0
\end{aligned}
$$

Now suppose that $P_{t+1, A_{t}} \leq P_{t A_{t}}$. By [Lattimore and Szepesvári, 2019, Theorem 26.5],

$$
\left\langle P_{t}-P_{t+1}, \hat{\ell}_{t}\right\rangle-D_{F_{t}}\left(P_{t+1}, P_{t}\right) \leq \frac{1}{2}\left\|\hat{\ell}_{t}\right\|_{\left(\nabla^{2} F_{t}(z)\right)^{-1}}^{2}
$$

where $z=\alpha P_{t}+(1-\alpha) P_{t+1}$ for some $\alpha \in[0,1]$. By definition $\nabla^{2} F_{t}(z)=\operatorname{diag}\left(\nabla^{2} f_{t}(z)\right) / \eta_{t}$ and since $\hat{\ell}_{t i}=0$ for $i \neq A_{t}$,

$$
\frac{1}{2}\left\|\hat{\ell}_{t}\right\|_{\nabla^{2} F_{t}(z)^{-1}}^{2}=\frac{\eta_{t} \hat{\ell}_{t A_{t}}^{2}}{2 \nabla^{2} f_{t}\left(z_{A_{t}}\right)} \leq \frac{\eta_{t} \hat{\ell}_{t A_{t}}^{2}}{2 \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}
$$

where we used the fact that $z_{A_{t}} \leq P_{t A_{t}}$ and that $\nabla^{2} f_{t}(p)$ is decreasing. The result follows by substituting the definition of $\hat{\ell}_{t A_{t}}$ and summing over $t \in[n]$.

Lemma 4.8. Let $\left(x_{t}\right)_{t=1}^{n}$ be a sequence with $x_{t} \in[0, B]$ for all $t$. Then

$$
\sum_{t=1}^{n} \frac{x_{t}}{\sqrt{1+\sum_{s=1}^{t-1} x_{s}}} \leq 4 \sqrt{1+\frac{1}{2} \sum_{t=1}^{n} x_{t}}+B
$$

The proof follows from a comparison to an integral and is given in the supplementary material.

Proof of Theorem 4.1. Using the result of Theorem 3.1, Lemma 4.7 and the assumption on the difference in the potentials, we have

$$
R_{n} \leq h_{2}(n)+\mathbb{E}\left[\frac{h_{1}}{\eta_{n+1}}+\sum_{t=1}^{n} \frac{\eta_{t}}{2} \frac{\ell_{t A_{t}}^{2}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}\right]
$$

As $\ell_{t A_{t}} \leq 1$, we can apply Lemma 4.8 with

$$
x_{t}=\frac{\ell_{t A_{t}}^{2}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)} \leq B
$$

It follows that $\eta_{t}=\eta_{0} / \sqrt{1+\sum_{s=1}^{t-1} x_{s}}$, and thus

$$
\begin{aligned}
& \sum_{t=1}^{n} \frac{\eta_{t}}{2} \frac{\ell_{t A_{t}}^{2}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)} \leq \\
& \quad 2 \eta_{0} \sqrt{1+\frac{1}{2} \sum_{t=1}^{n} \frac{\ell_{t A_{t}}^{2}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}}+\frac{\eta_{0}}{2} B .
\end{aligned}
$$

The first term in the last line is proportional to $1 / \eta_{n+1}$, therefore using the definition of $\eta_{t}$, Jensen's inequality and $\ell_{t A_{t}}^{2} \leq \ell_{t A_{t}}$, the regret can be bounded as

$$
\begin{aligned}
R_{n} \leq & h_{2}(n)+\frac{\eta_{0}}{2} B+\left(\frac{\sqrt{2} h_{1}}{\eta_{0}}+2 \eta_{0}\right) \\
& \times \sqrt{1+\frac{1}{2} \mathbb{E}\left[\sum_{t=1}^{n} \frac{\ell_{t A_{t}}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}\right]} .
\end{aligned}
$$

The first bound in the theorem follows from

$$
\begin{aligned}
& \mathbb{E}\left[\sum_{t=1}^{n} \frac{\ell_{t A_{t}}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}\right] \\
& \quad \leq B \mathbb{E}\left[\sum_{t=1}^{n}\left(\ell_{t A_{t}}-\ell_{t 1}+\ell_{t 1}\right)\right]=B R_{n}+B L_{n 1}
\end{aligned}
$$

and then from choosing $\eta_{0}=\sqrt{h_{1}} / 2^{1 / 4}$ and solving the resulting quadratic equation with respect to $R_{n}$.

For the second bound, we use $l_{t A_{t}} \leq 1$ and the definition of $C$, such that

$$
\mathbb{E}\left[\sum_{t=1}^{n} \frac{\ell_{t A_{t}}}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}\right] \leq \frac{C n}{2}
$$

To prove the corollaries, we need to bound $h_{1}, h_{2}(n), B$, and $C$ :

Lemma 4.9. The Hessian of the hybrid potential in Corollary 4.3 is monotonically decreasing, and for $n \geq 3$
$\frac{1}{p^{2} \nabla^{2} f_{t}(p)} \leq \sqrt{k} \log ^{1+q} n, \quad \mathbb{E}\left[\frac{1}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}\right] \leq 2 \sqrt{k}$,
Proof. For $p \in \operatorname{interior}\left(\Delta^{k-1}\right)$,

$$
\nabla^{2} f_{t}(p)=\frac{1}{2 p^{3 / 2}}+\frac{1}{p^{2} \sqrt{k} \log ^{1+q} \max \{3, t\}}
$$

is a decreasing function of $p$. It follows that for $n \geq 3$

$$
\frac{1}{p^{2} \nabla^{2} f_{t}(p)} \leq \sqrt{k} \log ^{1+q} n
$$

Moreover,

$$
\begin{aligned}
\sup _{t, P_{t} \in \mathcal{A}_{t}} \mathbb{E} & {\left[\frac{1}{P_{t A_{t}}^{2} \nabla^{2} f_{t}\left(P_{t A_{t}}\right)}\right] } \\
& \leq \sup _{t, P_{t} \in \Delta^{k-1}} \mathbb{E}\left[\frac{2}{\sqrt{P_{t A_{t}}}}\right]=2 \sqrt{k} .
\end{aligned}
$$

Lemma 4.10. Under the conditions of Corollary 4.3,

$$
\begin{aligned}
v_{n} \leq & \frac{\sqrt{k}}{\eta_{n+1}}\left(\frac{4}{3}+\frac{2}{q}\right)+\frac{5.5 k}{\eta_{0}} \sqrt{1+9 k^{3 / 2} \log ^{1+q}\left(9 k^{3 / 2}\right)} \\
& +\frac{3.7 \sqrt{k}}{\eta_{0}} \sqrt{1+3 \sqrt{k} \log ^{1+q} 3}+2 k \log n
\end{aligned}
$$

Proof. Due to the chopped simplex and the factorised potential, we have (recall the definition in Theorem 3.1 and bound the second sum with the integral, as shown in the supplementary material)

$$
\begin{aligned}
v_{n}= & \sum_{t=1}^{n} d_{t}\left(g_{t}+(t-1)\right) \\
& \leq \sum_{t=1}^{n} \frac{2 k}{t^{2}}\left(\frac{1}{\eta_{t}} \sup _{p \in[1 / t, 1]}\left|\nabla f_{t}(p)\right|+(t-1)\right) \\
& \leq \sum_{t=1}^{n} \frac{2 k}{t^{2}}\left(\frac{1}{\eta_{t}} \sup _{p \in[1 / t, 1]}\left|\nabla f_{t}(p)\right|\right)+2 k \log n .
\end{aligned}
$$

For $p \in[1 / t, 1]$ the gradient is bounded as

$$
\begin{aligned}
& \left|\nabla f_{t}(p)\right| \leq \frac{1}{\sqrt{p}}+\frac{1}{p \sqrt{k} \log ^{1+q} \max \{3, t\}} \\
& \quad \leq \sqrt{t}+\frac{t}{\sqrt{k} \log ^{1+q} \max \{3, t\}}
\end{aligned}
$$

Therefore, the corresponding sum in $v_{n}$ converges. By a straightforward calculation (see the supplementary material), the Hessian is bounded as in Lemma 4.9),

$$
\begin{aligned}
v_{n} \leq & \frac{\sqrt{k}}{\eta_{n+1}}\left(\frac{4}{3}+\frac{2}{q}\right)+\frac{5.5 k}{\eta_{0}} \sqrt{1+9 k^{3 / 2} \log ^{1+q}\left(9 k^{3 / 2}\right)} \\
& +\frac{3.7 \sqrt{k}}{\eta_{0}} \sqrt{1+3 \sqrt{k} \log ^{1+q} 3}+2 k \log n
\end{aligned}
$$

Lemma 4.11. Under the conditions of Corollary 4.3,

$$
\begin{aligned}
& \min _{p \in \mathcal{A}_{n+1}}\left(F_{n+1}(p)+\left\|p-e_{1}\right\|_{1} n\right)-F_{1}\left(P_{1}\right) \\
& \quad+\sum_{t=1}^{n}\left(F_{t}\left(P_{t+1}\right)-F_{t+1}\left(P_{t+1}\right)\right) \\
& \quad \leq \frac{\sqrt{k}}{\eta_{n+1}}\left(3+\frac{1}{q}\right)+k+\frac{\sqrt{k}}{3 \eta_{0}} \sqrt{1+3 \sqrt{k} \log ^{1+q} 3} .
\end{aligned}
$$

Proof. The potential is a mixture of the INF and the logbarrier parts, $F_{t}(p)=-\frac{2}{\eta_{t}} \sum_{i} \sqrt{p_{i}}-\frac{\alpha_{t}}{\eta_{t}} \sum_{i} \log p_{i}$ with $\alpha_{t}=1 /\left(\sqrt{k} \log ^{1+q} \max \{3, t\}\right)$.

To control the contribution of the INF term, first notice that the INF part of $F_{n+1}(p)$ is negative. Moreover,

$$
\left(-\frac{2}{\eta_{t}}+\frac{2}{\eta_{t+1}}\right) \sum_{i=1}^{k} \sqrt{P_{t+1, i}} \leq 2 \sqrt{k}\left(\frac{1}{\eta_{t+1}}-\frac{1}{\eta_{t}}\right)
$$

Summing with the INF part of $-F_{1}\left(P_{1}\right)$ and telescoping shows that it contributes at most $2 \sqrt{k} / \eta_{n+1}$ to the sum. For log-barrier, suppose $\alpha_{t} / \eta_{t} \leq \alpha_{t+1} / \eta_{t+1}$. Then

$$
\left(-\frac{\alpha_{t}}{\eta_{t}}+\frac{\alpha_{t+1}}{\eta_{t+1}}\right) \sum_{i=1}^{k} \log \left(P_{t+1, i}\right) \leq 0
$$

Now suppose that $\alpha_{t} / \eta_{t}>\alpha_{t+1} / \eta_{t+1}$. For $t \geq 3$, as $P_{t+1} \in \mathcal{A}_{t+1}$,

$$
\begin{aligned}
\left(-\frac{\alpha_{t}}{\eta_{t}}+\right. & \left.\frac{\alpha_{t+1}}{\eta_{t+1}}\right) \sum_{i=1}^{k} \log \left(P_{t+1, i}\right) \\
& \leq\left(\frac{\alpha_{t}}{\eta_{t}}-\frac{\alpha_{t+1}}{\eta_{t+1}}\right) k \log (t+1) \\
& \leq \frac{1}{\eta_{t}}\left(\frac{\log (t+1)}{\log ^{1+q}(t)}-\frac{1}{\log ^{q}(t)}\right) \sqrt{k} \\
& \leq \frac{\sqrt{k}}{\eta_{t} t \log ^{1+q} t}
\end{aligned}
$$

Summing over $t$ and noting that due to $\alpha_{1}=\alpha_{2}=\alpha_{3}$ the potential is unchanged,

$$
\begin{aligned}
& \sum_{t=1}^{n}\left(-\frac{\alpha_{t}}{\eta_{t}}+\frac{\alpha_{t+1}}{\eta_{t+1}}\right) \sum_{i=1}^{k} \log w_{t+1, i} \\
& \quad \leq \sum_{t=3}^{n} \frac{\sqrt{k}}{\eta_{t} t \log ^{1+q} t} \leq \frac{\sqrt{k}}{\eta_{n+1} q}+\frac{\sqrt{k}}{3 \eta_{3}}
\end{aligned}
$$

where the last inequality (shown in the supplementary material) essentially compares the sum to the integral of $1 /\left(t \log ^{1+q} t\right)$ and uses that $1 / \log ^{q} t \leq 1$ for $t \geq 3$. We can further bound $\eta_{3}$ as

$$
\frac{1}{\eta_{3}} \leq \frac{1}{\eta_{0}} \sqrt{1+3 \sqrt{k} \log ^{1+q} 3}
$$

by using the fact that the Hessian is bounded (see the proof of Lemma 4.9).

Finally, the log-barrier part of $-F_{1}\left(P_{1}\right)$ is negative. The log-barrier part of $F_{n+1}(p)$ is bounded by $\sqrt{k} / \eta_{n+1}$ as $p \in \mathcal{A}_{n+1}$. Thus,

$$
\begin{aligned}
& \min _{p \in \mathcal{A}_{n+1}} F_{n+1}(p)+n\left\|p-e_{1}\right\|_{1} \\
& \quad \leq \frac{\sqrt{k}}{\eta_{n+1}}+\min _{p \in \mathcal{A}_{n+1}} n\left\|p-e_{1}\right\|_{1}=\frac{\sqrt{k}}{\eta_{n+1}}+\frac{k n}{n+1} .
\end{aligned}
$$

Combining the three bounds and using that $k n /(n+1) \leq$ $k$ concludes the proof.

Proof of Corollary 4.3. From Lemma 4.9, Lemma 4.10 and Lemma 4.11, we find

$$
\begin{aligned}
& B=\sqrt{k} \log ^{1+q} n, C=2 \sqrt{k}, h_{1}=\sqrt{k}\left(\frac{13}{3}+\frac{3}{q}\right), \\
& h_{2}(n)=2 k \log n+\frac{5.5 \cdot k}{\eta_{0}} \sqrt{1+9 k^{3 / 2} \log ^{1+q}\left(9 k^{3 / 2}\right)} \\
& \quad+\frac{4.1}{\eta_{0}} \sqrt{k} \sqrt{1+3 \sqrt{k} \log ^{1+q} 3}+k .
\end{aligned}
$$

Now applying Theorem 4.1 with $\eta_{0}=k^{1 / 4} \sqrt{\frac{13}{3 \sqrt{2}}+\frac{3}{\sqrt{2 q}}}$ completes the proof. Note that in the big-O notation, we only kept the leading terms that grow with $n$.

Proof of Corollary 4.4. Starting from the end of the previous proof, choosing $q=1$ and upper-bounding the numerical coefficients, we obtain the corollary.

## 5 VARIANCE OF THE REGRET

The expected regret is just one measure of the performance of an algorithm. Algorithms with small expected regret may suffer from a large variance. Since the adversarial model is often motivated on the grounds of providing robustness, it would be unfortunate if proposed algorithms suffered from high variance. Recently, however, it was shown that the variance of Exp3 without exploration is quadratic in the horizon [Lattimore and Szepesvári, 2019, §11], and a similar result holds for Thompson sampling in a Bayesian setting [Bubeck and Sellke, 2019]. Here we generalise these arguments to prove quadratic variance of the regret for a class of algorithms based on FTRL with importanceweighted loss estimators. This is the worst possible result for bandits with bounded losses. The class of policies covered by our theorem includes INF and Exp3, but not FTRL with the log barrier. To keep things simple we restrict ourselves to algorithms of the form

$$
P_{t}=\underset{p \in \Delta^{k-1}}{\arg \min }\left\langle p, \hat{L}_{t-1}\right\rangle+\frac{1}{\eta_{n}} \sum_{i=1}^{k} f\left(p_{i}\right),
$$

where $f$ is convex and $\left(\eta_{n}\right)_{n=1}^{\infty}$ is a sequence of learning rates. Note that this corresponds to a sequence of algorithms, each with a fixed learning rate.
Assumption 5.1. The number of actions is $k=2$ and $f$ is Legendre with $(0,1) \subseteq \operatorname{dom}(f)$ and $0 \in \partial \operatorname{dom}(f)$.

The assumption on the potential is satisfied by all standard potentials for bandits on the probability
simplex, including those in Table 1. It allows us to write $P_{t}$ in a simple form. Let $g(p)=f(p)+f(1-p)$, which is convex and Legendre with $\operatorname{dom}(g)=(0,1)$. Given $x \geq 0$,

$$
\underset{p \in[0,1]}{\arg \min }(p x+g(p))=\nabla g^{*}(-x),
$$

where we used the fact that for Legendre functions the gradient is invertible and $(\nabla g)^{-1}=\nabla g^{*}$. That $g$ is Legendre with $\operatorname{dom}(g)=(0,1)$ also ensures that $\nabla g^{*}$ is nondecreasing and $\lim _{x \rightarrow-\infty} \nabla g^{*}(x)=0$ and $\lim _{x \rightarrow \infty} \nabla g^{*}(x)=1$. By symmetry, we also have $\nabla g^{*}(0)=1 / 2$. The point is that by the definition of FTRL, $P_{t 1}=\nabla g^{*}\left(\eta_{n}\left(\hat{L}_{t-1,2}-\hat{L}_{t-1,1}\right)\right)$.
Theorem 5.2. Assume $\limsup _{n \rightarrow \infty} n \nabla g^{*}\left(-a n \eta_{n}\right)<$ $\infty$ for all $a>0$. Then for all sufficiently large $n$ there exists a bandit for which $\mathbb{P}\left(\hat{R}_{n} \geq n / 4\right) \geq c$, where $c>0$ is a constant that depends on the algorithm, but not the horizon.
Corollary 5.3. Under the same conditions as Theorem 5.2 the variance of the regret is $\operatorname{Var}\left[\hat{R}_{n}\right]=\Omega\left(n^{2}\right)$.

Examples Suppose $\eta_{n}=a n^{-1 / 2}$ for some $a>0$. Then the conditions of the theorem are satisfied when $f$ is the negentropy. In this case $\nabla g^{*}$ is the sigmoid function and the corresponding algorithm is just Exp3. When $f(p)=-2 \sqrt{p}$ and $x \leq 0$, then

$$
\nabla g^{*}(x)=\frac{1}{2}\left(1-\sqrt{1+\frac{4\left(2 \sqrt{1+x^{2}}-2-x^{2}\right)}{x^{4}}}\right),
$$

which satisfies $\lim \sup _{n \rightarrow \infty} \nabla g^{*}(-a \sqrt{n}) n=1 / a^{2}$. In this sense $1 / 2$-Tsallis entropy with $\eta_{n}=\Theta\left(n^{-1 / 2}\right)$ just barely satisfies the conditions. The consequence is that the minimax optimal INF policy proposed by Audibert and Bubeck [2009] has quadratic variance. The log barrier does not satisfy the conditions and we speculate it is more stable.

Proof of Theorem 5.2. Assume for simplicity that 4 is a factor of $n$. Let $\alpha_{n} \in[0,1 / 2]$ be a constant to be tuned subsequently and consider a bandit defined by

$$
\ell_{t 1}=\left\{\begin{array}{ll}
\alpha_{n} & \text { if } t \leq n / 2 \\
0 & \text { otherwise } .
\end{array} \quad \ell_{t 2}= \begin{cases}0 & \text { if } t \leq n / 2 \\
1 & \text { otherwise } .\end{cases}\right.
$$

Clearly the first arm is optimal. Let $c_{1}>0$ be a constant such that for all sufficiently large $n$ it holds that $\nabla g^{*}\left(-n \eta_{n}\right) \leq c_{1} / n$, which is guaranteed to exist by the assumptions in the theorem. Then define events $F_{t}=\cap_{s=n / 2+1}^{t}\left\{A_{s}=2, P_{s 1} \leq c_{1} / n\right\}$. On the event $F_{n}$ the random regret satisfies

$$
\begin{equation*}
\hat{R}_{n} \geq \frac{n}{2}-\frac{\alpha_{n} n}{2} \geq \frac{n}{4} . \tag{4}
\end{equation*}
$$

The theorem follows by proving that $\mathbb{P}\left(F_{n}\right) \geq c$ for all sufficiently large $n$ and constant $c>0$. The idea is to show that the estimated loss for the optimal arm after the first $n / 2$ rounds is large enough that the algorithm never plays the optimal arm in the second half of the game with constant probability.

First half dynamics The choice of $\alpha_{n}$ determines the dynamics of the interaction between the algorithm and environment in the first $n / 2$ rounds. Before the main proof we establish some facts about this. Let $\alpha \in[0,1 / 2]$ and define $\left(p_{s}(\alpha)\right)_{s=0}^{n}$ inductively by $p_{0}(\alpha)=1 / 2$ and

$$
p_{s+1}(\alpha)=\nabla g^{*}\left(-\eta_{n} \sum_{u=0}^{s} \frac{\alpha}{p_{u}(\alpha)}\right)
$$

which is chosen so that $P_{t+1,1}=p_{s}(\alpha)$ whenever $t+1 \leq n / 2$ and $T_{1}(t)=s$. Here we used the fact that $\hat{L}_{t 2}=0$ for $t \leq n / 2$, which follows from the definition of the bandit. Let $Q_{s}(\alpha)=\sum_{u=0}^{s-1} \alpha / p_{u}(\alpha)$. Clearly $Q_{2}(1 / 2)>0$ and $Q_{s}(0)=0$ for all $s$. Furthermore, $Q_{s}(\alpha)$ is increasing in both $\alpha$ and $s$ and continuous in $\alpha$. Therefore there exists an $\alpha_{\circ} \in(0,1 / 2)$ such that $Q_{2}(1 / 2) \geq Q_{3}\left(\alpha_{\circ}\right)$. Now suppose that $Q_{s}(1 / 2) \geq$ $Q_{s+1}\left(\alpha_{\circ}\right)$. Using the fact that $\nabla g^{*}$ is increasing,

$$
\begin{aligned}
& Q_{s+1}(1 / 2)=Q_{s}(1 / 2)+\frac{1}{2} \nabla g^{*}\left(-\eta_{n} Q_{s}(1 / 2)\right)^{-1} \\
& \geq Q_{s+1}\left(\alpha_{\circ}\right)+\alpha_{\circ} \nabla g^{*}\left(-\eta_{n} Q_{s+1}\left(\alpha_{\circ}\right)\right)^{-1}=Q_{s+2}\left(\alpha_{\circ}\right),
\end{aligned}
$$

which by induction means that $Q_{s}(1 / 2) \geq Q_{s+1}\left(\alpha_{\circ}\right)$ for all $s \geq 2$. Notice that $\hat{L}_{t 1}=Q_{s}\left(\alpha_{n}\right)$ when $T_{1}(t-1)=s$.

Second half dynamics Define threshold $\lambda_{n}$ by

$$
\lambda_{n}=n+n^{2} /\left(2\left(n-c_{1}\right)\right) \leq 2 n
$$

where the latter inequality holds for all sufficiently large $n$. Let $E$ be the event $E=\left\{\hat{L}_{n / 2,1} \geq \lambda_{n}\right\}$. We claim that $\mathbb{P}\left(F_{n} \mid E\right) \geq \exp \left(-c_{1} / 2\right)$. Suppose that $t>n / 2$ and $E \cap F_{t}$ occurs. Then
$\hat{L}_{t 2}=\sum_{s=n / 2+1}^{t} \frac{1}{P_{s 2}} \leq \sum_{s=n / 2+1}^{t} \frac{1}{1-c_{1} / n} \leq \frac{n^{2}}{2\left(n-c_{1}\right)}$
where the first inequality follows from the definition of $F_{t}$. Therefore, since $\hat{L}_{t 1} \geq \hat{L}_{n / 2,1} \geq \lambda_{n}$,

$$
\begin{align*}
P_{t+1,1} & =\nabla g^{*}\left(\eta_{n}\left(\hat{L}_{t 2}-\hat{L}_{t 1}\right)\right) \\
& \leq \nabla g^{*}\left(\eta_{n}\left(\frac{n^{2}}{2\left(n-c_{1}\right)}-\lambda_{n}\right)\right) \leq \frac{c_{1}}{n} \tag{5}
\end{align*}
$$

Hence $\mathbb{P}\left(F_{t+1} \mid F_{t}, E\right) \geq 1-c_{1} / n$. Noting that Eq. (5) implies that $P_{n / 2+1,1} \leq c_{1} / n$ shows that $E \subseteq F_{n / 2+1}$ and hence by induction

$$
\begin{equation*}
\mathbb{P}\left(F_{n} \mid E\right) \geq\left(1-\frac{c_{1}}{n}\right)^{n / 2} \geq \exp \left(-c_{1} / 2\right) \tag{6}
\end{equation*}
$$

Lower bounding $\mathbb{P}(E) \quad$ By Eqs. (4) and (6) it suffices to prove that $\mathbb{P}(E)$ is larger than a constant for sufficiently large $n$. Let $s=\min \left\{u: Q_{u}(1 / 2) \geq \lambda_{n}\right\}$, which by our assumptions on $\nabla g^{*}$ for sufficiently large $n$ is at least $s>2$ and at most $s \leq n / 2$. Then $Q_{s}\left(\alpha_{\circ}\right) \leq$ $Q_{s-1}(1 / 2)<\lambda_{n} \leq Q_{s}(1 / 2)$. By the intermediate value theorem and the continuity of $\alpha \mapsto Q_{s}(\alpha)$ we may choose $\alpha_{n} \in\left(\alpha_{\circ}, 1 / 2\right]$ such that $Q_{s}\left(\alpha_{n}\right)=\lambda_{n}$. Now introduce a sequence of independent geometric random variables $\left(G_{u}\right)_{u=0}^{s}$ with $G_{u} \in\{1,2, \ldots\}$ and $\mathbb{E}\left[G_{u}\right]=$ $1 / p_{u}(\alpha)$. Then by construction,

$$
\begin{equation*}
\mathbb{P}\left(T_{1}(n / 2) \geq s\right)=\mathbb{P}\left(\sum_{u=0}^{s-1} G_{u} \leq \frac{n}{2}\right) \tag{7}
\end{equation*}
$$

You should think of $G_{u}$ as the number of rounds before the algorithm plays action 1 for the $u$ th time. Let

$$
\kappa=\min \left\{m: \sum_{u=0}^{s-m-1} \frac{1}{p_{u}\left(\alpha_{n}\right)} \leq \frac{n}{8}\right\}
$$

Then either $\sum_{u=0}^{s-\kappa-1} 1 / p_{u}\left(\alpha_{n}\right) \leq n / 16$ in which case $1 / p_{s-\kappa}\left(\alpha_{n}\right) \geq n / 16$ or $\sum_{u=0}^{s-\kappa-1} 1 / p_{u}\left(\alpha_{n}\right) \geq n / 16$. Then there exists a constant $c_{2} \geq 0$ such that for sufficiently large $n$,

$$
\begin{aligned}
p_{s-\kappa}\left(\alpha_{n}\right) & =\nabla g^{*}\left(-\eta_{n} Q_{s-\kappa-1}\left(\alpha_{n}\right)\right) \\
& =\nabla g^{*}\left(-\eta_{n} \sum_{u=0}^{s-\kappa-1} \frac{\alpha_{n}}{p_{u}\left(\alpha_{n}\right)}\right) \\
& \leq \nabla g^{*}\left(-\frac{\alpha_{\circ} n \eta_{n}}{16}\right) \leq \frac{c_{2}}{n} .
\end{aligned}
$$

Combining the two cases and choosing $c_{2} \geq 16$ guarantees that $p_{s-\kappa}\left(\alpha_{n}\right) \leq c_{2} / n$ for sufficiently large $n$. Using the fact that $s \mapsto p_{s}\left(\alpha_{n}\right)$ is decreasing,

$$
2 n \geq \lambda_{n}=\sum_{u=0}^{s-1} \frac{\alpha_{n}}{p_{u}\left(\alpha_{n}\right)} \geq \sum_{u=s-\kappa}^{s-1} \frac{\alpha_{\circ} n}{c_{2}}=\frac{\kappa \alpha_{\circ} n}{c_{2}}
$$

Rearranging shows that $\kappa$ is less than a constant that is independent of $n$. By Markov's inequality

$$
\begin{aligned}
& \mathbb{P}\left(\sum_{u=0}^{s-\kappa-1} G_{u} \geq \frac{n}{4}\right) \\
& \quad \leq \mathbb{P}\left(\sum_{u=0}^{s-\kappa-1} G_{u} \geq 2 \sum_{u=0}^{s-\kappa-1} \frac{1}{p_{u}\left(\alpha_{n}\right)}\right) \leq \frac{1}{2}
\end{aligned}
$$

Hence

$$
\begin{equation*}
\mathbb{P}\left(\sum_{u=0}^{s-\kappa-1} G_{u}<\frac{n}{4}\right) \geq \frac{1}{2} \tag{8}
\end{equation*}
$$

Furthermore,

$$
\begin{aligned}
\frac{\alpha_{\circ}}{p_{s-1}\left(\alpha_{n}\right)} & \leq \frac{\alpha_{n}}{p_{s-1}\left(\alpha_{n}\right)} \leq \sum_{u=0}^{s-1} \frac{\alpha_{n}}{p_{u}\left(\alpha_{n}\right)} \\
& =Q_{s}\left(\alpha_{n}\right)=\lambda_{n} \leq 2 n
\end{aligned}
$$

Therefore, using again that $s \mapsto p_{s}(\alpha)$ is decreasing,

$$
\begin{aligned}
& \mathbb{P}\left(\sum_{u=s-\kappa}^{s-1} G_{u} \leq \frac{n}{4}\right) \\
& \quad \geq\binom{ n / 4}{\kappa} p_{s-1}\left(\alpha_{n}\right)^{\kappa}\left(1-p_{s-\kappa}\left(\alpha_{n}\right)\right)^{n / 4-\kappa} \\
& \quad \geq\binom{ n / 4}{\kappa}\left(\frac{\alpha_{\circ}}{2 n}\right)^{\kappa}\left(1-\frac{c_{2}}{n}\right)^{n / 4-\kappa}
\end{aligned}
$$

which for sufficiently large $n$ is larger than a strictly positive constant and the result follows by combining the above with Eqs. (7) and (8).

Remark 5.4. We believe the result continues to hold for adaptive learning rates under the assumption that $\lim \sup _{t \rightarrow \infty} t \nabla g^{*}\left(-a t \eta_{t}\right)<\infty$ for all $a>0$. The proof becomes significantly more delicate, however.

## 6 LINEARLY SEPARABLE BANDITS

In this section we consider the case where the adversary chooses an infinite sequence of loss vectors $\left(\ell_{t}\right)_{t=1}^{\infty}$. The main objective is to prove logarithmic (or better) regret under the following assumption.
Assumption 6.1. There is a linear separation between the optimal and suboptimal arms:

$$
\Delta_{i}=\liminf _{n \rightarrow \infty}\left(L_{n i}-L_{n 1}\right) / n>0 \quad \text { for all } i>1
$$

Note that if $\left(\ell_{t}\right)_{t=1}^{\infty}$ are independent and identically distributed random vectors, then the above holds almost surely whenever there is a unique optimal arm. We provide two results in this setting. The first generalises a known result from stochastic bandits that there exist algorithms for which the asymptotic random regret grows arbitrarily slowly almost surely [Cowan and Katehakis, 2015].
Theorem 6.2. For any nondecreasing function $f: \mathbb{N} \rightarrow$ $\mathbb{N}$ with $\lim _{n \rightarrow \infty} f(n)=\infty$ there exists an algorithm such that $\lim \sup _{n \rightarrow \infty} \hat{R}_{n} / f(n)<\infty$ almost surely.

The algorithm realising the bound in Theorem 6.2 explores uniformly at random on a set $E$ for which $\limsup _{n \rightarrow \infty}|E \cap[n]| / f(n) \leq 1$ almost surely. The reader is warned that the constants hidden by the asymptotics are potentially quite enormous.

Of course this result says nothing about the expected regret, which must be logarithmic for consistent algorithms [Lai and Robbins, 1985]. The following theorem improves on a result by Seldin and Slivkins [2014] by a factor of $\log (n) / \log \log (n)$.
Theorem 6.3. There exists an algorithm such that for any adversarial bandit $R_{n}=O(\sqrt{k n})$. Furthermore, under Assumption 6.1 it holds that

$$
\limsup _{n \rightarrow \infty} \frac{R_{n}}{\log (n)^{2} \log \log (n)}<\infty
$$

The algorithm is INF with enough forced exploration that the loss estimators are guaranteed to be sufficiently accurate to detect a linear separation. The proofs of Theorems 6.2 and 6.3 use standard concentration results and are given in the supplementary material.

## 7 OPEN QUESTIONS

Despite the relatively long history and extensive research, many open questions exist about $k$-armed adversarial bandits. Perhaps the most exciting question is the existence/nature of a genuinely instance-optimal algorithm. The work by Zimmert and Seldin [2019] suggests the possibility of an algorithm for which $R_{n}=$ $O(\sqrt{k n})$ and $R_{n}=O\left(\sum_{i: \Delta_{i}>0} \log (n) / \Delta_{i}\right)$, where $\Delta_{i}=\frac{1}{n} \sum_{t=1}^{n}\left(\ell_{t i}-\ell_{t 1}\right)$ is the empirical gap between the arms. In fact, one could hope for a little more. For stochastic Bernoulli bandits with means $\left(\theta_{i}\right)_{i=1}^{k}$, the KL-UCB algorithm by Cappé et al. [2013] satisfies $R_{n}=O\left(\sum_{i: \Delta_{i}>0} \Delta_{i} \log (n) / d\left(\theta_{i}, \theta^{*}\right)\right)$ where $d\left(\theta_{i}, \theta_{1}\right)$ is the relative entropy between Bernoulli distributions with bias $\theta_{i}$ and $\theta_{1}$ respectively. We are not aware of a lower bound proving that such a result is not possible for adversarial bandits with $\theta_{i}=\frac{1}{n} \sum_{t=1}^{n} \ell_{t i}$. At present it is not clear whether or not our modified algorithm from Corollary 4.3 retains the logarithmic regret in the stochastic setting, both because we use an adaptive learning rate and a hybrid potential. Finally, it is known that sub-exponential tail bounds are incompatible with logarithmic regret in the stochastic setting [Audibert et al., 2009], but by appropriately tuning the confidence intervals it is straightforward to prove the variance is linear in $n$, which is optimal. Missing is an adaptation of INF that enjoys (a) minimax regret, (b) logarithmic regret in the stochastic setting and (c) linear variance.

## Acknowledgements

This work was supported by the Gatsby Charitable Foundation. The authors thank Haipeng Luo for spotting an error in the earlier version of the manuscript.

## References

J. Abernethy and A. Rakhlin. Beating the adaptive bandit with high probability. In 2009 Information Theory and Applications Workshop, pages 280-289, 2009. 2
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 17th International Conference on Algorithmic Learning Theory, ALT, pages 229-243, Berlin, Heidelberg, 2006. Springer-Verlag. 2
J.-Y. Audibert and S. Bubeck. Minimax policies for adversarial and stochastic bandits. In Proceedings of Conference on Learning Theory (COLT), pages 217226, 2009. 1, 7
J.-Y. Audibert and S. Bubeck. Best arm identification in multi-armed bandits. In Proceedings of Conference on Learning Theory (COLT), 2010. 1, 2
J.-Y. Audibert, R. Munos, and Cs. Szepesvári. Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. Theoretical Computer Science, 410(19):1876-1902, 2009. 9
P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. Gambling in a rigged casino: The adversarial multiarmed bandit problem. In Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on, pages 322-331. IEEE, 1995. 2
S. Bubeck and N. Cesa-Bianchi. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. Foundations and Trends in Machine Learning. Now Publishers Incorporated, 2012. 1
S. Bubeck and M. Sellke. First-order regret analysis of thompson sampling. arXiv preprint arXiv:1902.00681, 2019. 7
S. Bubeck and A. Slivkins. The best of both worlds: Stochastic and adversarial bandits. In COLT, pages 42.1-42.23, 2012. 1
S. Bubeck, M. Cohen, and Y. Li. Sparsity, variance and curvature in multi-armed bandits. In F. Janoos, M. Mohri, and K. Sridharan, editors, Proceedings of Algorithmic Learning Theory, volume 83 of Proceedings of Machine Learning Research, pages 111-127. PMLR, 07-09 Apr 2018. 2
O. Cappé, A. Garivier, O. Maillard, R. Munos, and G. Stoltz. Kullback-Leibler upper confidence bounds for optimal sequential allocation. The Annals of Statistics, 41(3):1516-1541, 2013. 9
W. Cowan and M. N. Katehakis. Asymptotic behavior of minimal-exploration allocation policies: Almost sure, arbitrarily slow growing regret. arXiv preprint arXiv:1505.02865, 2015. 1, 9
S. Gerchinovitz and T. Lattimore. Refined lower bounds for adversarial bandits. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, Advances in Neural Information Processing Systems 29, NIPS, pages 1198-1206. Curran Associates, Inc., 2016. 2
E. Hazan. Introduction to online convex optimization. Foundations and Trends ${ }^{\circledR}$ in Optimization, 2(3-4): 157-325, 2016. 2
E. Hazan and S. Kale. A simple multi-armed bandit algorithm with optimal variation-bounded regret. In S. M. Kakade and U. von Luxburg, editors, Proceedings of the 24th Annual Conference on Learning Theory, volume 19 of Proceedings of Machine Learning Research, pages 817-820. PMLR, 2011. 2
P. Joulani, A. György, and Cs. Szepesvári. A modular analysis of adaptive (non-)convex optimization: Optimism, composite objectives, and variational bounds. In $A L T$, 2017. 3
T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. Advances in applied mathematics, 6(1):4-22, 1985. 9
T. Lattimore and Cs. Szepesvári. Bandit Algorithms. Cambridge University Press (preprint), 2019. 1, 2, 5, 7
G. Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, Advances in Neural Information Processing Systems 28, NIPS, pages 3168-3176. Curran Associates, Inc., 2015a. 2
G. Neu. First-order regret bounds for combinatorial semi-bandits. In P. Grünwald, E. Hazan, and S. Kale, editors, Proceedings of The 28th Conference on Learning Theory, volume 40 of Proceedings of Machine Learning Research, pages 1360-1375, Paris, France, 03-06 Jul 2015b. PMLR. 2
Y. Seldin and G. Lugosi. An improved parametrization and analysis of the EXP3++ algorithm for stochastic and adversarial bandits. In S. Kale and O. Shamir, editors, Proceedings of the 2017 Conference on Learning Theory, volume 65 of Proceedings of Machine Learning Research, pages 1743-1759, Amsterdam, Netherlands, 07-10 Jul 2017. PMLR. 1
Y. Seldin and A. Slivkins. One practical algorithm for both stochastic and adversarial bandits. In E. P. Xing and T. Jebara, editors, Proceedings of the 31st International Conference on Machine Learning, volume 32 of Proceedings of Machine Learning Research, pages 1287-1295, Bejing, China, 22-24 Jun 2014. PMLR. 1, 9
S. Shalev-Shwartz. Online learning: Theory, algorithms, and applications. PhD thesis, The Hebrew University of Jerusalem, 2007. 2
C-Y. Wei and H. Luo. More adaptive algorithms for adversarial bandits. In S. Bubeck, V. Perchet, and P. Rigollet, editors, Proceedings of the 31st Conference On Learning Theory, volume 75 of Proceedings of Machine Learning Research, pages 1263-1291. PMLR, 06-09 Jul 2018. 2, 4
J. Zimmert and Y. Seldin. An optimal algorithm for stochastic and adversarial bandits. In Kamalika Chaudhuri and Masashi Sugiyama, editors, Proceedings of Machine Learning Research, volume 89 of Proceedings of Machine Learning Research, pages 467-475. PMLR, 16-18 Apr 2019. 1, 4, 9
J. Zimmert, H. Luo, and C-Y. Wei. Beating stochastic and adversarial semi-bandits optimally and simultaneously. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, Proceedings of the 36th International Conference on Machine Learning, volume 97 of Proceedings of Machine Learning Research, pages 7683-7692. PMLR, 09-15 Jun 2019. 2, 4

