
Adapting to game trees in zero-sum imperfect information games

Côme Fiegel¹ Pierre Ménard² Tadashi Kozuno³ Rémi Munos⁴ Vianney Perchet^{1,5} Michal Valko⁴

Abstract

Imperfect information games (IIG) are games in which each player only partially observes the current game state. We study how to learn ε -optimal strategies in a zero-sum IIG through self-play with *trajectory feedback*. We give a problem-independent lower bound $\tilde{O}(H(A_{\mathcal{X}} + B_{\mathcal{Y}})/\varepsilon^2)$ on the required number of realizations to learn these strategies with high probability, where H is the length of the game, $A_{\mathcal{X}}$ and $B_{\mathcal{Y}}$ are the total number of actions for the two players. We also propose two Follow the Regularized leader (FTRL) algorithms for this setting: **Balanced FTRL** which matches this lower bound, but requires the knowledge of the information set structure beforehand to define the regularization; and **Adaptive FTRL** which needs $\tilde{O}(H^2(A_{\mathcal{X}} + B_{\mathcal{Y}})/\varepsilon^2)$ realizations without this requirement by progressively adapting the regularization to the observations.

1. Introduction

In imperfect information games (IIG), players, upon taking an action, may only have access to *partial information about the true current game state*. This type of games allows for the modelling of complex strategic behavior such as bluffing (Koller & Pfeffer, 1995). In this work, in particular, we study extensive-form two-players zero-sum IIGs.

The *extensive-form* of a game describes, which player is playing and which actions are available, sequentially and depending on the previous moves. It is typically represented by a tree of depth H , where nodes are *states* that are controlled by one of the players that decides, based on each player action, the next game states among its children.

Players can be uncertain about the true game state upon

¹CREST, ENSAE, IP Paris, Palaiseau, France ²ENS Lyon, Lyon, France ³Omron Sinic X, Tokyo, Japan ⁴Deepmind, Paris, France ⁵CRITEO AI Lab, Paris, France. Correspondence to: Côme Fiegel <come.fiegel@normalesup.org>.

playing. We therefore assume that the set of states controlled by the player is partitioned into *information sets*, where an information set is a set of states that are indistinguishable to this player. As commonly assumed in game theory, we suppose *perfect recall* (Kuhn, 1950): players remember their previous moves, which implies that the space of information sets has a *tree structure*.

We focus more specifically on *zero-sum games* (what the max-player gains is the opposite of the other, min-player) and aim at devising algorithms that learn ε -optimal strategies (von Neumann, 1928). For this purpose, we actually develop a unilateral algorithm that dictates to a single player how to play over multiple realizations of the game called *episodes*, each of the same length H , in order to minimize the *regret*—the difference between the cumulative gain and the gain of the best fixed strategy—against any sequence of moves of the adversary. If both players follows such an algorithm, we can then prove that their average strategy becomes more precise over time. We especially want to construct algorithms with a regret that scales as slowly as possible with respect to the different game parameters.

For any information set $x \in \mathcal{X}$, we define by A_x the number of available actions at x , as two indistinguishable states must have the same number of actions. The total number of actions of the max-player is then defined by $A_{\mathcal{X}} = \sum_{x \in \mathcal{X}} A_x$. Similarly, we let \mathcal{Y} be the collection of information sets of the min-player, B_y the number of actions available at $y \in \mathcal{Y}$ and $B_{\mathcal{Y}} = \sum_{y \in \mathcal{Y}} B_y$ the total number of actions. Meanwhile X and Y denote the number of information sets for the max and min-players.

Related work When the structure of the game, transition probabilities and reward function are known beforehand, several methods exist to approximate the optimal strategy. A common approach is to reformulate the problem as a linear program that can be solved efficiently (Romanovsky, 1962; von Stengel, 1996; Koller et al., 1996). Some methods see it as a saddle-point problem and rely on first-order optimization methods (Hoda et al., 2010; Kroer et al., 2015; 2018; 2020; Munos et al., 2020; Lee et al., 2021). Other common approaches try to locally minimize the regret at each information set, e.g., *counterfactual regret minimization* (Zinkevich et al., 2007; Tammelin, 2014; Burch et al., 2019).

The main practical weakness of the above approaches is their prohibitive cost as the size of the game grows. Indeed, the computations are mostly done on the whole information set trees, which implies a time complexity of at least $O(X + Y)$ at each step at best (Lanctot et al., 2009). All the recent successes in empirically solving large IIGs actually avoid the direct use of these full information algorithms (Moravčík et al., 2017; Brown & Sandholm, 2018; Schmid et al., 2021; Bakhtin et al., 2022; Polat et al., 2022). In addition, depending on the practical case, it may not be possible to know the state transitions, or even to start the game at an arbitrary state.

We therefore treat the settings when the algorithms can only access trajectory feedback of the game: they only have access to the observations of the player. For this case, the outcome sampling MCFR (Lanctot et al., 2009; Schmid et al., 2018; Farina et al., 2020) feeds estimates of the counterfactual regret to the CFR algorithm, which is by design a full-information feedback algorithm. Precisely, outcome-sampling MCFR uses a uniform sampling of the actions in each information set, which is actually sub-optimal when the information sets of the players are unbalanced, as the information sets in smaller sub-trees are then observed more than information sets of bigger sub-trees. For this reason, Farina et al. (2020) instead propose to build an estimator of the counterfactual regret according to a balanced policy that, roughly speaking, samples actions proportionally to the size of the associated sub-trees. While this balanced policy allows for a better sample complexity, it also requires to know the game structure beforehand.

Closer to the adversarial bandits, Farina et al. (2021b) propose to use the online mirror descent (OMD) algorithm along with a dilatation of the Shannon entropy on all information sets acting as a regularizer (Hoda et al., 2010). However, their approach uses importance sampling (Auer et al., 2003) to estimate the losses, whose variance is too large to allow the computation of -optimal strategies with high probability. Kozuno et al. (2021) solved this issue by using a biased estimate called implicit exploration (IX, Kocák et al., 2014; Neu, 2015; Lattimore & Szepesvári, 2020), and proved that their IXOMD algorithm achieves a sample complexity of order $O(H^2(XA_X + YB_Y))$ with high probability. Unfortunately, it does not match the lower bound, see Table 1.

More recently, Bai et al. (2022) propose to weight differently each of the information set in the regularizer using the aforementioned balanced policy. They obtain a sample com

¹The exact definition of the balanced policy varies among different works, but the main idea remains the same.

²For algorithms with a probability at least ϵ of a correct output, the symbol \tilde{O} hides dependencies logarithmic in $H; A_X; B_Y$; and n .

plexity of $O(H^3(A_X + B_Y))$ that scales linearly with the total number of actions, but again, at the extra price of requiring the knowledge of the game structure beforehand. Hence, we raise the following question: "What is the optimal rate for learning -optimal strategies through self-play? Is it possible to learn these optimal strategies with a complexity scaling linearly with the total number of actions without the prior knowledge of the information set tree structure? In this work we answer both questions

Contributions We make the following contributions:

We prove a lower bound $O(H(A_X + B_Y))$ on the number of necessary plays required to compute an -optimal profile with only a trajectory feedback. Compared to the one stated by Bai et al. (2022), it applies to algorithms that are only correct with a high probability, and enjoys an extra factor if the size of the action set is allowed to vary with the information set.

We propose the **Balanced FTRL** algorithm that matches this lower bound, up to logarithm factors, using the knowledge of the structure of the information sets. It uses the Follow The Regularized Leader (FTRL) algorithm instead of OMD and follows the idea of balancing, but using instead a concept of balanced transition. In addition to the use of Shannon entropy, it also allows the use of the Tsallis entropy, in order to get a better sample complexity at the price of increased computation time.

We then propose the **Adaptive FTRL** algorithm which matches this lower bound up to an extra logarithmic factors, without using the knowledge of the structure. Instead of using the balanced transition, it estimates the actual transitions of each episode.

We provide a practical implementation of **Balanced FTRL** and **Adaptive FTRL** through Algorithm 3 of Appendix F. This implementation is computationally efficient, as the policy is only updated on the information sets visited in the previous episode.

Finally, we provide experiments that compare the performances of these two algorithms with OMD and Balanced OMD. We observe that, despite having different theoretical guarantees, the algorithms all seem to have comparable performances in practice.

All rates are summarized in Table 1.

2. Setting

Consider an episodic, finite, two-player, zero-sum IIG $(S; X; Y; f_A(x); g_{2X}; f_B(y); g_{2Y}; H; p; r)$, which consists of the following components (Kuhn, 1950; Littman, 1994):

- A finite state space S and two information set spaces (partitions of S) X of size X and Y of size Y for the max- and min-player respectively.

Algorithm	Sample complexity	Structure-free
MCCFR (Karina et al., 2020; Bai et al., 2022)	$\mathcal{O}(H^4(A_X + B_Y)^2)$	7
IXOMD (Kozuno et al., 2021)	$\mathcal{O}(H^2(XA_X + YB_Y)^2)$	3
Balanced OMD (Bai et al., 2022)	$\mathcal{O}(H^3(A_X + B_Y)^2)$	7
Balanced FTRL (this paper)	$\mathcal{O}(H(A_X + B_Y)^2)$	7
Adaptive FTRL (this paper)	$\mathcal{O}(H^2(A_X + B_Y)^2)$	3
Lower bound (this paper)	$\mathcal{O}(H(A_X + B_Y)^2)$	

Table 1. Sample complexity for episodic, finite, two-player, zero-sum imperfect information games. Structure-free algorithm designates an algorithm that does not need to know the structure of the information set spaces in advance. The symbols $\mathcal{O}(H^k(A_X + B_Y)^2)$ and $\mathcal{O}(H^k(XA_X + YB_Y)^2)$ indicate dependencies logarithmic in $A_X + B_Y$ and $XA_X + YB_Y$ respectively. Note that for algorithms that only work with fixed action sets of size A_X and B_Y we have $A_X = |A_X|$ and $B_Y = |B_Y|$.

- For each $x \in X$ and $y \in Y$, finite action spaces $A(x)$ of size A_X and $B(y)$ of size B_Y .
- The length $H \in \mathbb{N}$ of the game.
- Initial state distribution $p_0 \in \Delta(S)$ and a state-transition probability kernel $(p_h)_{h \in [H]}$ with $p_h : S \times A(X) \times B(Y) \rightarrow \Delta(S)$ for each $h \in [H]$.
- A reward function $(r_h)_{h \in [H]}$ with $r_h : S \times A(X) \times B(Y) \rightarrow [0, 1]$.

At step t , the max- and min-player observe their information sets s_t and y_t . Given the information, the max- and min-player choose and execute actions $a_t \in A(x_t)$ and $b_t \in B(y_t)$. As a result, the current state transitions to a next state $s_{t+1} \sim p_h(s_t, a_t, b_t)$, and the max- and min-player receive rewards $r_h(s_t, a_t, b_t)$ and $r_h(s_t, a_t, b_t)$ respectively. This is repeated until time step H , after which the episode finishes.

Perfect-recall As explained in the introduction, we assume perfect-recall. Formally, this means that for the max-player and for each information set $x \in X$ there exists a unique $h \in [H]$ and history (x_1, a_1, \dots, x_h) such that $x_h = x$. Specifically, we write x^0 if x^0 is part of the history that leads to x . With this assumption, both X and Y can be partitioned into different subsets $(X_h)_{h \in [H]}$ and $(Y_h)_{h \in [H]}$, X_h and Y_h being the sets of possible information sets at time step h for respectively the max and min-player.

Learning procedure We assume there are T episodes of the same game, and both players are able to progressively adapt their respective policies π and σ after each episode based on their previous observations.

Realization plan and loss Given the perfect recall assumption, we recursively define (von Stengel, 1996) the realization plan $\pi_{1:h} := (\pi_{1:h})_{h \in [H]}$ by, given (x_1, a_1, \dots, x_h) the unique history up to h ,

$$\pi_{1:h}(x_h; a_h) := \prod_{h^0=1}^h p_{h^0}(a_{h^0} | x_{h^0})$$

For convenience we let $A(X_h) := \{f(x_h; a_h) : x_h \in X_h \text{ and } a_h \in A(x_h)\}$ be the total action set for the max-player at depth h , and use an analogous notation $B(Y_h)$ for the min-player. The unions of these sets are denoted by $A(X) := \bigcup_{h \in [H]} A(X_h)$; $B(Y) := \bigcup_{h \in [H]} B(Y_h)$ with their respective sizes A_X and B_Y , as defined in the introduction.

We also define the adversarial transitions $p_{1:h} := (p_{1:h})_{h \in [H]}$ and losses $\ell_h := (\ell_h)_{h \in [H]}$ by:

$$p_{1:h}(x_h) := p_0(x_1) \prod_{h^0=1}^{h-1} p_{h^0}(x_{h^0+1} | x_{h^0}, a_{h^0});$$

$$\ell_h(x_h; a_h) := p_{1:h}(x_h) (1 - r_h(x_h; a_h));$$

Policies The perfect-recall assumption allows us to represent a policy of the max-player as a sequence $(\pi_h)_{h \in [H]}$ such that for all $x_h \in X_h$, $\pi_h(\cdot | x_h)$ is an element of $\Delta(A(x_h))$. A policy σ of the min-player can be defined similarly. We let π_{\max} and π_{\min} be the sets of the max- and min-player's policies.

where $p_h(x_{h+1} | x_h; a_h)$ is the probability to transition to the information state x_{h+1} from the pair information set-action $(x_h; a_h)$ when the opponent policy is fixed to σ , and $r_h(x_h; a_h)$ is the average reward obtained in this case. Note that these two quantities require the perfect recall assumption in order to be well defined, and combine the randomness of both state transitions and min-player's policy

Episode unfolding Given the two players' policies π and σ , an episode of the game proceeds as follows: an initial

³In many games with tree structure, p_h does not depend on the step h ; but we keep the dependence on h as it makes the results more general at no extra cost in the analysis.

⁴As in Kozuno et al. (2021); Bai et al. (2022), we assume that the players can not change their policies within an episode, as the tree-like structure limits the interest of such adaptation.

The analog quantities q_1, p_1, r_h and \cdot may also be defined for the min-player, replacing r_h by r_h in the definition of the loss.

Regret and ϵ -Nash-equilibrium We define the expected rewards (of the max-player) $V^t := E \sum_{h=1}^H r_h$ for a pair (π^t, σ^t) of policies and, after T episodes, the regrets of the max and min players,

$$R_{\max}^T := \max_{\pi^t} \sum_{t=1}^T (V^{\pi^t, \sigma^t} - V^{\pi^t, \sigma^t});$$

$$R_{\min}^T := \max_{\sigma^t} \sum_{t=1}^T (V^{\pi^t, \sigma^t} - V^{\pi^t, \sigma^t});$$

There is a very clear link between minimizing the regret of both players and computing optimal policies, or ϵ -Nash-equilibrium (NE). In the particular case of perfect-recall, Theorem 1 by [Kozuno et al. \(2021\)](#) states the following (see also [Cesa-Bianchi & Lugosi 2006](#); [Zinkevich et al. 2007](#)).

Theorem 2.1. From $(\pi^t, \sigma^t)_{t \in [T]}$ define the time-averaged policy $(\bar{\pi}, \bar{\sigma})$ (of the realization plan, see Appendix E), then $(\bar{\pi}, \bar{\sigma})$ is ϵ -optimal with

$$\epsilon = \frac{R_{\max}^T + R_{\min}^T}{T} = \epsilon_T;$$

Conversion to online linear regret minimization Defining a scalar product between the realization plans and the losses,

$$\langle \pi_{1:t}, \ell_{1:t} \rangle := \sum_{h=1}^H \sum_{a_h} \pi_{1:t}(x_h; a_h) \ell_h(x_h; a_h);$$

the chain rule on conditional probabilities implies that the max-player regret can then be rewritten, with $\ell_{1:t}$, as

$$R_{\max}^T = \max_{\pi^t} \sum_{t=1}^T \langle \pi^t, \ell_{1:t} \rangle - \langle \pi^t, \ell_{1:t} \rangle;$$

This converts the problem to a linear regret problem, where the constraints set (the set of realization plans) is a polytope.

For simplicity, in the following sections we focus on the max-player although the exact same ideas apply to the min-player, who aims at minimizing R_{\min}^T .

3. Lower bounds

In this section we state lower bounds for learning an approximate NE. Compared to [Bai et al. \(2022\)](#), our lower bounds cover high-probability guarantees instead of in expectation, and are stated for two settings, whether or not the size of the action set is fixed. As shown in Section 4, they are tight up to logarithmic factors, which justifies the existence of two different rates for these settings.

Theorem 3.1. (informal, exact statement in Appendix C)

Regret: For any algorithm that controls the max-player, there exists a game such that, if the number of action per information set is allowed to vary, with probability at least $1 - \delta$ the algorithm suffers a regret of

$$\sqrt{H A_X T \log(1/\delta)};$$

Otherwise if the number of the actions is fixed, then the regret is

$$\sqrt{A_X T \log(1/\delta)};$$

Sample complexity: For any algorithm that controls both players, there exists a game such that, if the number of action is allowed to vary, it needs at least

$$H(A_X + B_Y) \log(1/\delta) = \Omega^2$$

episodes to output an optimal strategy with a probability at least $1 - \delta$.

If the number of actions is fixed, then it instead needs

$$(A_X + B_Y) \log(1/\delta) = \Omega^2 \text{ episodes.}$$

These lower bounds are obtained by a reduction to the k -armed bandits. The structure depends on the setting: when the number of actions is allowed to vary, the learner must only choose between k actions on the first information set, and then gets k times the same reward depending on this first choice. When the number of actions is fixed, the learner must take H actions successively (among k proposed at each information set) and only receive a reward after the last one depending on all of its choices; hence with $k = A^H$ possibilities.

4. Optimal rate with the knowledge of the tree structure

We first study the case where the information set structure is known by the agent beforehand.

Loss estimation The agent only observes the realizations of the games. This suggests that the loss vector should be estimated using only the information on the visited states and the associated rewards. Therefore, we define, for all possible actions, the unbiased loss estimate

$$b_h^t(x_h; a_h) := \frac{1_{\{x_h^t = x_h; a_h^t = a_h\}}}{1_{1:h}(x_h; a_h)} (1 - r_h^t);$$

where x_h^t and r_h^t are respectively the information set, and reward of the max-player at time step t .

While this estimator works well in expectation, its variance is too large to get acceptable bounds with high probability. We replace it by the following (biased) estimator with an implicit exploration (IX) parameter ϵ_h that depends on the chosen state (Kak et al., 2014):

$$\hat{\epsilon}_h(x_h; a_h) := \frac{\mathbb{1}_{x_h^t = x_h; a_h^t = a_h} g}{\sum_{1:h}(x_h; a_h) + \epsilon_h(x_h)} (1 - r_h^t)$$

The associated cumulative loss is given as $\sum_{k=1}^t \epsilon_k$.

Balanced transitions We follow the idea of Bai et al. (2022) and define balanced transition kernel $p^?$ on X . First, we denote by $A_x = \sum_{x^0 \dots x^h} A_{x^0}$ the total number of actions that follow information set x in the whole sub-tree. The transition kernel $p^?$ is then defined such as, at each step, the probability of transitioning to the next suitable information set is proportional to this value. More precisely, $p_{x_1}^?(x_1) = A_{x_1} = A_x$ for all roots $x_1 \in X_1$, and for any $h \in [H-1]$,

$$p_{x_{h+1}|x_h; a_h}^? := \frac{A_{x_{h+1}}}{\sum_{x_{h+1} \in X_{h+1}(x_h; a_h)} A_{x_{h+1}}},$$

where $X_{h+1}(x_h; a_h)$ denotes all the possible information sets of depth $h+1$ that follow $(x_h; a_h)$.

This definition differs from the balanced transition p^h in the appendix of Bai et al. (2022) for two reasons:

- It is defined globally for all depths h , using directly the whole tree instead of a truncated version up to each depth.
- It uses the reachable number of action instead of the reachable number of states, as the number of action at each set may not be constant.

Policy update (U1) For any transition kernel p on X and max-player policy π , the vector $p_{1:h} : 1:h$ defined by $p_{1:h} : 1:h(x_h; a_h) := p_{1:h}(x_h) p_{1:h}(x_h; a_h)$ is a probability distribution over the information set-action X_h at depth h . Along the lines of FTRL algorithms, we define the policy update of **Balanced FTRL** as

$$x_h^t = \underset{x_h \in X_h}{\operatorname{argmin}} \sum_{1:h} \sum_{a_h} \epsilon_{1:h}^t(x_h; a_h) + \sum_{h=1}^t \sum_{x_h \in X_h} p_{1:h}^?(x_h) \sum_{a_h} \pi_{1:h}^t(x_h; a_h) \quad (\text{U1})$$

where $(\epsilon_h)_{h \in [H]}$ are regularizers of $A(X_h)$ and $(\epsilon_h)_{h \in [H]}$ is a sequence of learning rates. While the ideal choice for regret minimization would be the actual adversarial transitions p^t (replacing $p^?$) these transitions are unknown by the player. The balanced transitions instead give acceptable bounds against any adversarial transition. **Balanced FTRL** implements this idea, using either Tsallis entropy (Tsallis, 1988) or Shannon entropy as regularizer. Similarly, the IX parameters ϵ_h are also adjusted for each state with $\epsilon_h^?(x_h) = \epsilon_h = p_{1:h}^?(x_h)$.

Algorithm 1 **Balanced FTRL-Tsallis/Shannon**

1: Input:
 Tree-like structure $\mathcal{A}(X)$
 Sequence $(\epsilon_h)_{h \in [H]}$ of learning rates
 Sequence $(\epsilon_h)_{h \in [H]}$ of IX parameters
 Regularizers $\phi_h : X \rightarrow \mathbb{R}^+$:
 Tsallis: $\phi_h(w_h) = \frac{1}{q} \sum_{x_h \in X_h} w_h(x_h; a_h)^q$
 Shannon: $\phi_h(w_h) = \sum_{x_h \in X_h} w_h(x_h; a_h) \log(w_h(x_h; a_h))$

2: Initialize:
 For all depth h and $x_h \in X_h$:
 $\epsilon_h^?(x_h) = \epsilon_h = p_{1:h}^?(x_h)$

3: For $t = 1$ to T
 Compute update (U1):
 $x_h^t = \underset{x_h \in X_h}{\operatorname{argmin}} \sum_{1:h} \sum_{a_h} \epsilon_{1:h}^t(x_h; a_h) + \sum_{h=1}^t \sum_{x_h \in X_h} p_{1:h}^?(x_h) \sum_{a_h} \pi_{1:h}^t(x_h; a_h)$
 For $h = 1$ to H
 Observe information set x_h^t
 Execute $a_h^t(x_h^t)$ and receive reward r_h^t
 Compute loss:
 $\epsilon_h^?(x_h^t; a_h^t) = \epsilon_{1:h}^t(x_h^t; a_h^t) + \epsilon_h^?(x_h^t)$

Time complexity **Balanced FTRL** first needs to initially compute the balanced transitions, which can be done recursively from the bottom of the information set tree with a time complexity of $O(A_x)$.

Then, with **Balanced FTRL-Tsallis**, each update (U1) can be done by solving a convex programming over X_h . While this implies a computation with a polynomial cost in A_x , it is still quite expensive when A_x gets large. On the other hand, we will see in Section 5 that the update of **Balanced FTRL-Shannon** has an equivalent formulation with the dilated Shannon entropy. Using Algorithm 3 of Appendix F, each update can then be computed with a time complexity of $O(HA)$, where A is an upperbound on the local number of actions any information set can have.

Theoretical analysis We now state the main theoretical results for **Balanced FTRL** with proof in Appendix E.

Theorem 4.1. Let $q \in (0; 1)$ and $\epsilon = \log(3A_x)$. Fix the IX parameters to $\epsilon_h = \frac{\epsilon}{H} = \frac{\epsilon}{2A_x T}$ for all $h \in [H]$.

Using **Balanced FTRL-Tsallis**, $q \in (0; 1)$ and constant learning rates $\epsilon_h = \frac{\epsilon}{2q(1-q)T} = \frac{\epsilon}{2q(1-q)H}$, with probability at least $1 - \frac{1}{T}$,

$$R_{\max}^T \leq \frac{\epsilon}{2(q(1-q))} + 3 \frac{\epsilon}{2} \frac{\epsilon}{HA_x T}$$

Using **Balanced FTRL-Shannon** and constant learning rates $\epsilon_h = \frac{\epsilon}{2H \log(A_x)} = \frac{\epsilon}{2A_x T}$, with probability at least $1 - \frac{1}{T}$,

$$R_{\max}^T \leq \frac{\epsilon}{2 \log(A_x)} + 3 \frac{\epsilon}{2} \frac{\epsilon}{HA_x T}$$

Remark 4.2 The best theoretical rate is obtained with Dilated entropy and policy update (U2) For a given Tsallis entropy with $q = 1=2$, similarly to Abernethy et al. (2015). Even though Shannon entropy has an extra $\log(A_X)$ factor in the (expected) first term of the bound as in the multi-armed bandits setting, both entropies give a rate $O(\sqrt{H \log(A_X + B_Y)})$ with high-probability. Indeed, the expected term is dominated in both cases by the term from the high-probability analysis.

Using Theorem 2.1, we conclude that the lower bound $\Theta(\sqrt{H(A_X + B_Y)})$ on the sample complexity is matched up to logarithmic factors when both players use Balanced FTRL. In particular, Balanced FTRL improves by a factor H over the rate of Balanced OMD by Bai et al. (2022), see Table 1.

Corollary 4.3. For any $\epsilon > 0$ and $\beta \in (0, 1)$, if both player run either Balanced FTRL-Tsallis/-Shannon with the parameters specified in Theorem 4.1, (η, γ) is an ϵ -optimal profile with probability at least $1 - \beta$ as long as

$$T \geq O\left(\frac{H(A_X + B_Y) \log(1/\beta)}{\epsilon^2}\right)$$

The average profile can be computed along the update of the current profile, see Kozuno et al. (2021), at the price of a local update with $O(A_X + B_Y)$ time complexity.

Fixed action set size When the size of the action set is fixed, the number of information sets having a given X in its history will increase at least exponentially as the depth increases. This can be exploited by Balanced FTRL with suitable learning rates and IX parameters that both decrease exponentially with the depth in order to remove the dependency in the upper bounds, again nearly matching the lower bounds $\Theta(\sqrt{H(A_X + B_Y)})$ given by Theorem 3.1 for this setting. The results, proved in Appendix E, are stated for Shannon entropy:

Theorem 4.4. Assume that the size of the action set is constant equal to A and let $\beta \in (0, 1)$, $\eta = \frac{\log(3A)}{2A}$. Then, Balanced FTRL-Shannon with $\eta_h = \frac{\eta}{A^{H-h}}$ and $\gamma_h = \frac{\eta}{2A^{H-h} \log(A)}$ yields

$$R_{\max}^T \leq 5 \sqrt{\log(A)} + 8 \sqrt{\frac{1}{A}}$$

As in Corollary 4.3, this allows to compute an ϵ -optimal profile as long as long as

$$T \geq O\left(\frac{H(A_X + B_Y) \log(1/\beta)}{\epsilon^2}\right)$$

5. Near-optimal rate without the knowledge of the tree structure

In this section we drop the assumption on the knowledge of the tree structure of the information states.

vector η of learning rates η^k , we define the weighted dilated Shannon divergence (Kroer et al., 2015) between two policies π^1, π^2 as

$$D_{\eta}(\pi^1, \pi^2) := \sum_{h=1}^H \sum_{(x_h, a_h) \in \mathcal{X}_h} \frac{\eta_h(x_h, a_h)}{\pi_h^1(x_h)} \log \frac{\pi_h^2(a_h | x_h)}{\pi_h^1(a_h | x_h)}$$

We are now interested in the updates

$$\eta^t = \operatorname{argmin}_{\eta} D_{\eta}(\pi^1, \pi^2) + D_{\eta}(\eta^{t-1}, \eta^t) \quad (U2)$$

for $\eta^0 \in \mathcal{H}$ a base policy fixed over all iterations. As shown by Algorithm 3 of Appendix F, these updates enjoy an efficient computation with $O(HA)$ time complexity if the vectors η^t of learning rates only change at the visited information sets $\{x_1^t, \dots, x_H^t\}$.

Connections with the balanced algorithm While there is no general equivalence between the Shannon entropy and the dilated Shannon entropy, the updates of Balanced FTRL-Shannon may be put in the form of updates (U2). Precisely, using the fact that π_h^t is a transition kernel over \mathcal{X}_h , the updates (U1) of Balanced FTRL can be rewritten for all $t \in [T]$ as (see Proposition F.2)

$$\eta^t = \operatorname{argmin}_{\eta} D_{\eta}(\pi^1, \pi^2) + D_{\eta}(\eta^{t-1}, \eta^t);$$

for some $\beta \in (0, 1)$ and vector η^0 of learning rates computed with $O(A_X)$ time complexity.

Adversarial transitions estimation The balanced adversarial transitions $\pi_{1:h}^t(x_h)$ used in the above learning rates $\eta^t(x_h)$ are impossible to compute if the tree structure is unknown at the start of the game. In analogy to the loss estimation, we could instead use the estimates of the actual adversarial transitions $\pi_h^t(x_h)$. For any $(x_h, a_h) \in \mathcal{X}_h$, an unbiased estimate of $\pi_h^t(x_h)$ is indeed given by $\hat{\pi}_{1:h}^t(x_h, a_h) := \mathbb{I}_{f_{x_h=x_h^t, a_h=a_h^t}} g^t = \pi_{1:h}^t(x_h, a_h)$.

Once more, we shall use IX estimators to get bounds with high probability and not just in expectation. However, Balanced FTRL used the IX parameter $\pi_h^t(x_h) = \pi_{1:h}^t(x_h)$ which also can not be computed for the same reason. Our idea is to again replace these transitions by the same estimated transitions we are currently trying to define. This leads to the following recursive definitions,

$$\begin{aligned} \pi_h^t(x_h, a_h) &:= \pi_{1:h}^t(x_h, a_h) + \beta^t \pi_{1:h}^{t-1}(x_h, a_h); \\ \hat{\pi}_{1:h}^t(x_h, a_h) &:= \frac{\mathbb{I}_{f_{x_h=x_h^t, a_h=a_h^t}} g^t}{\pi_{1:h}^t(x_h, a_h) + \pi_{1:h}^{t-1}(x_h, a_h)}, \end{aligned}$$

where $\beta^t := \prod_{k=1}^t \beta^k$ are the cumulative estimated transitions and $\pi_{1:h}^t(x_h, a_h)$ the IX parameter that now also depends on the action a_h and on the time t .

Algorithm 2 Adaptive FTRL

1: Input:
 Learning rate η and IX base parameter β
 Base policy π^0 (uniform by default)

2: For $t = 1$ to T :
 For all depth h and $x_h \in X_h$:

$$\hat{p}_{1:h}^t(x_h) = \min_{x_{h^0} \in X_{h^0}} \frac{1}{\beta} \log \left(\sum_{a_h \in A(x_h)} \exp(\beta \sum_{s \in S} \pi_s^0(x_h, a_h)) \right)$$

 For all $a_h \in A(x_h)$:

$$\hat{p}_{1:h}^t(x_h; a_h) = \frac{1}{\beta} \log \left(\sum_{s \in S} \pi_s^0(x_h, a_h) \right)$$

 Compute update (U2) $\hat{\pi}_h^t = \arg \min_{\pi_h} \sum_{a_h \in A(x_h)} \pi_h(a_h) \left(\sum_{s \in S} \pi_s^0(x_h, a_h) + D(\pi_h; \pi^0) \right)$
 For $h = 1$ to H :
 Observe information set x_h^t
 Execute $a_h^t = \hat{\pi}_h^t(x_h^t)$ and receive reward r_h^t

$$\hat{q}_h^t(x_h^t; a_h^t) = (1 - \beta) r_h^t + \beta \left(\hat{p}_{1:h}^t(x_h^t; a_h^t) + \hat{p}_{1:h}^t(x_h^t; a_h^t) \right)$$

$$\hat{p}_{1:h}^t(x_h^t; a_h^t) = \frac{1}{\beta} \log \left(\sum_{s \in S} \pi_s^0(x_h^t, a_h^t) + \hat{q}_h^t(x_h^t; a_h^t) \right)$$

ditional $H \log(T) \log(1 + T)$ factor on the sample complexity compared to the optimal rate. This result is proved in Appendix G.

Theorem 5.1. Let $\beta = \log(3A_X T)$ and $\eta = 1 + \log_2(1 + T)$. For any $\epsilon \in (0, \frac{1}{2})$, using Adaptive FTRL with $\beta = \frac{1}{2\epsilon} \log(3A_X T)$, $\eta = \frac{1}{2\epsilon} \log(1 + T)$ and π^0 the uniform policy yields with a probability of at least $1 - \epsilon$,

$$R_{\max}^T \leq 6H \frac{1}{\epsilon} \log(3A_X T)$$

As in Corollary 4.3, this leads to the computation of an ϵ -optimal profile as long as

$$T \geq O \left(H^2 (A_X + B_Y) \log(1 + T) \log(T) \right)^{\frac{1}{\epsilon^2}}$$

The rate of Adaptive FTRL improves the one of XOMD—up till now the algorithm with the best rate without the knowledge of the structure—by a factor at least $\frac{1}{\epsilon} \log(A_X Y)$.

Adaptive learning rates To replace the learning rate $\eta_h^t(x_h)$ associated to each information set $x_h \in X_h$, we thus use the cumulative estimated transitions. As at each round t for each action $a_h \in A(x_h)$, the quantity $\hat{p}_{1:h}^t(x_h; a_h)$ is an estimator of $\pi_{1:h}^t(x_h)$, we define the estimator $\hat{p}_{1:h}^t(x_h)$, as the average of these estimations

$$\hat{p}_{1:h}^t(x_h) := \frac{1}{|A(x_h)|} \sum_{a_h \in A(x_h)} \hat{p}_{1:h}^t(x_h; a_h) = A_{x_h} :$$

This estimator is then used to define a learning rate $\eta_h^t(x_h)$, mimicking the action of $\pi_h^t(x_h)$ in $\hat{p}_{1:h}^t(x_h)$,

$$\eta_h^t(x_h) := \min_{x_{h^0} \in X_{h^0}} \frac{1}{\beta} \log \left(\sum_{a_h \in A(x_h)} \exp(\beta \sum_{s \in S} \pi_s^0(x_{h^0}, a_h)) \right)$$

For technical reasons, these learning rates have to be increasing along a trajectory, hence the \max operator. Note that this definition at round t does not depend on whether the yet unexplored information states are considered in this min operator, as these verify $\pi_{1:h^0}^t(x_{h^0}^0) = 0$.

Adaptive algorithm Adaptive FTRL is based on update (U2) with the above vectors $\{\hat{\pi}_h^t\}_{t \in [T]}$ of learning rates and IX estimators, which allow it is naturally adaptive to the initially unknown tree structure; see Algorithm 2 for a detailed description of Adaptive FTRL.

Time complexity Unlike the Balanced FTRL algorithm, there is no initialization cost. Algorithm 3, in Appendix F, gives an efficient way of computing each update (U2) with a $O(HA)$ time complexity. Moreover, as β only increases on visited information states, the vectors $\hat{\pi}_h^t$ and $\hat{p}_{1:h}^t$ can be updated in-place with the same $O(HA)$ time complexity.

Theoretical results We show high-probability sample complexity bounds for Adaptive FTRL, but with an ad-

ditional $H^2 (A_X + B_Y) \log(1 + T) \log(T) = \frac{1}{\epsilon^2}$:
 Parameter tuning Adaptive FTRL needs the knowledge of the A_X constant but only in order to tune the learning rate and the IX parameter. Fortunately, the doubling trick can be applied to A_X similarly to how it would be applied to β , which means that, up to an additional constant factor, this knowledge is not required for the theoretical guarantees.

6. Experiments

In this section we provide preliminary experiments for Balanced FTRL and Adaptive FTRL on simple games. The code used to generate these experiments is available at <https://github.com/anon17893/IIG-tree-adaptation>.

Games We compare the algorithms on the following three standard benchmark tabular games: Kuhn poker (Kuhn, 1950), Leduc poker (Southey et al., 2005) and liars dice. We use the implementation of these games available in the OpenSpiel library (Lanctot et al., 2019).

Implementation details We implemented IXOMD (Kozuno et al., 2021) Balanced OMD (Bai et al., 2022), Balanced FTRL-Shannon and Adaptive FTRL. We also implemented Tweaked Adaptive FTRL, a slight modification of Adaptive FTRL, in which the learning rates and IX parameters decrease depending on $\beta^t = (1 + \beta^t)^{-1}$ instead of $\beta^t = (1 + \beta^t)$. This version gets better practical performances, at the price of not having tight theoretical guarantees.

These algorithms strongly rely on the learning rates for their performances. However, the learning rates obtained with the theoretical analysis of the above algorithms are only adjusted with respect to the worst possible player games

Figure 1. Performances of various algorithms with respect to the number of episodes for the time-averaged problem of Theorem 2.1. The results are given in terms of the exploitability $\max_{\pi} \min_{\pi'} V^{\pi, \pi'}$, with the rewards scaled between 0 and 1. The total number of actions is always the same for both players, with $A_X = 12$ for Kuhn poker, $A_X = 1092$ for Leduc poker, and $A_X = 24570$ for Liars dice.

of a given size. With such rates, the practical performances of **Adaptive FTRL** directly estimates the transition kernels induced by the opponent instead of relying on a balanced transition kernel. We thus tune the rates separately for each algorithm and each game, using a (logarithmic) grid search on the global learning rates, while the base η parameter was taken as $\eta = 20$ of this global learning rate.

Results Figure 1 shows the results of these experiments. We notice that, despite having very different theoretical guarantees, all algorithms seem to have very similar practical performances once tuned, except for **Adaptive FTRL** being a little worse and its modified version **Tweaked Adaptive FTRL** being slightly better on average. We believe that the behaviour of **Adaptive FTRL** can be explained by its learning rates of order $\frac{1}{\sqrt{t}}$, which is very large at the beginning leading to a fast convergence in the early phase. But the learning rates then need some time to decrease enough to allow **Adaptive FTRL** to make more progress⁵. **Tweaked Adaptive FTRL** stays adaptive but avoids this issue with a less steep learning rate decrease and a smaller initial learning rate.

7. Conclusion

We presented two algorithms **Balanced FTRL** is problem-independent optimal for learning the optimal problem. It requires to know the structure of the information set spaces in advance in order to compute a balanced transition kernel used to weight the regularizer. **Adaptive FTRL** does not require the knowledge of the information set space structure while being near problem-independent optimal. Indeed,

⁵Note that this effect is responsible for an additional $\frac{1}{\log(T)}$ factor in the regret analysis

Our results bring the following interesting future directions:

Optimal rate for Adaptive FTRL Is it possible to save the extra H dependence over the horizon in the sample complexity of **Adaptive FTRL**?

Last iterate convergence Can we obtain theoretical guarantees for the current problem instead of the average problem used in Theorem 2.1, in the same way as Lee et al. (2021)? Indeed, the computation of the time-averaged problem is not straightforward outside of the current tabular setting (Heinrich et al., 2015; Brown et al., 2019).

Avoiding importance sampling In large games, the importance exploration estimates tend to have a very high variance. This may be an issue if we want to combine the algorithms with function approximation (McAleer et al., 2022). Therefore, an interesting future direction would be to design algorithms that do not rely on neither the importance sampling nor the knowledge of the information set structure, but still get sharp sample complexity guarantees

Acknowledgements

P. Ménard acknowledges the Chaire SeqALO (ANR-20-CHIA-0020-01). Vianney Perchet acknowledges support from the French National Research Agency (ANR) under grant number ANR-19-CE23-0026 as well as from the grant "Investissements d'Avenir" (LabEx Ecodec/ANR-11-LABX-0047).

References

- Abernethy, J. D., Lee, C., and Tewari, A. Fighting bandits with a new kind of smoothness. *Neural Information Processing Systems* 2015. URL <https://proceedings.neurips.cc/paper/2015/file/5caf41d62364d5b41a893adc1a9dd5d4-Paper.pdf>.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. The Nonstochastic Multiarmed Bandit Problem. *SIAM Journal on Computing* 32(1):48–77, January 2003. ISSN 0097-5397. doi: 10.1137/S0097539701398375.
- Bai, Y., Jin, C., and Yu, T. Near-optimal reinforcement learning with self-play. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.) *Advances in Neural Information Processing Systems*, volume 33, pp. 2159–2170. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/172ef5a94b4dd0aa120c6878fc29f70c-Paper.pdf>.
- Bai, Y., Jin, C., Mei, S., and Yu, T. Near-optimal learning of extensive-form games with imperfect information. In *International Conference on Machine Learning* 2022.
- Bakhtin, A., Brown, N., Dinan, E., Farina, G., Flaherty, C., Fried, D., Goff, A., Gray, J., Hu, H., Jacob, A. P., Komeili, M., Konath, K., Kwon, M., Lerer, A., Lewis, M., Miller, A. H., Mitts, S., Renduchintala, A., Roller, S., Rowe, D., Shi, W., Spisak, J., Wei, A., Wu, D., Zhang, H., and Zijlstra, M. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science* 378(6624):1067–1074, 2022. doi: 10.1126/science.ade9097. URL <https://www.science.org/doi/abs/10.1126/science.ade9097>.
- Brown, N. and Sandholm, T. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science* 359(6374):418–424, 2018. doi: 10.1126/science.aao1733. URL <https://www.science.org/doi/abs/10.1126/science.aao1733>.
- Brown, N., Lerer, A., Gross, S., and Sandholm, T. Deep counterfactual regret minimization. *International Conference on Machine Learning* 2019. URL <https://proceedings.mlr.press/v97/brown19b.html>.
- Burch, N., Moravčík, M., and Schmid, M. Revisiting CFR+ and Alternating Updates. *Journal of Artificial Intelligence Research* 64:429–443, 2019.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, 2006. ISBN 978-0-521-84108-5. doi: 10.1017/CBO9780511546921.
- Daskalakis, C., Foster, D. J., and Golowich, N. Independent policy gradient methods for competitive reinforcement learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.) *Advances in Neural Information Processing Systems*, volume 33, pp. 5527–5540. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/3b2acfe2e38102074656ed938abf4ac3-Paper.pdf>.
- Farina, G., Kroer, C., and Sandholm, T. Regret circuits: Composability of regret minimizers. In Chaudhuri, K. and Salakhutdinov, R. (eds.) *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1863–1872. PMLR, 2019. URL <http://proceedings.mlr.press/v97/farina19b.html>.
- Farina, G., Kroer, C., and Sandholm, T. Stochastic Regret Minimization in Extensive-Form Games. *International Conference on Machine Learning* 2020.
- Farina, G., Kroer, C., and Sandholm, T. Faster Game Solving via Predictive Blackwell Approachability: Connecting Regret Matching and Mirror Descent. *AAI Conference on Artificial Intelligence* 2021a. URL <https://arxiv.org/abs/2007.14358>.
- Farina, G., Kroer, C., and Sandholm, T. Bandit Linear Optimization for Sequential Decision Making and Extensive-Form Games. In *AAI Conference on Artificial Intelligence* 2021b.
- Gilpin, A., Peña, J., and Sandholm, T. First-order algorithm with $O(\ln(1/\epsilon))$ convergence for ϵ -equilibrium in two-person zero-sum games. *Mathematical Programming* 133, 2012. URL <https://www.cs.cmu.edu/~sandholm/restart.MathProg12.pdf>.
- Gordon, G. J. No-regret Algorithms for Online Convex Programs. In *Advances in Neural Information Processing Systems* 2007.
- Hart, S. and Mas-Colell, A. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica* 68(5): 1127–1150, 2000.
- Heinrich, J., Lanctot, M., and Silver, D. Fictitious self-play in extensive-form games. *International conference on machine learning* 2015.
- Hoda, S., Gilpin, A., Peña, J., and Sandholm, T. Smoothing Techniques for Computing Nash Equilibria of Sequential Games. *Mathematics of Operations Research* 2010. URL <https://kilthub.cmu.edu/ndownloader/files/12101699>.
- Johanson, M., Bard, N., Lanctot, M., Gibson, R., and Bowling, M. Efficient nash equilibrium approximation through

- monte carlo counterfactual regret minimization. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2 AAMAS '12*, pp. 837–846, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 0981738125.
- Kocák, T., Neu, G., Valko, M., and Munos, R. Efficient learning by implicit exploration in bandit problems with side observations. *Neural Information Processing Systems*, 2014.
- Koller, D. and Pfeffer, A. Generating and solving imperfect information games. *International Joint Conference on Artificial Intelligence*, 1995.
- Koller, D., Megiddo, N., and Von Stengel, B. Efficient Computation of Equilibria for Extensive Two-Person Games. *Games and Economic Behavior*, 14(2):247–259, 1996.
- Kozuno, T., Ménard, P., Munos, R., and Valko, M. Learning in two-player zero-sum partially observable Markov games with perfect recall. *Neural Information Processing Systems*, 2021.
- Kroer, C., Waugh, K., Kilian-Karzan, F., and Sandholm, T. Faster first-order methods for extensive-form game solving. In *Economics and Computation*, 2015. ISBN 978-1-4503-3410-5. doi: 10.1145/2764468.2764476.
- Kroer, C., Farina, G., and Sandholm, T. Solving large sequential games with the excessive gap technique. *Neural Information Processing Systems*, 2018.
- Kroer, C., Waugh, K., Kilian-Karzan, F., and Sandholm, T. Faster algorithms for extensive-form game solving via improved smoothing functions. *Mathematical Programming*, 179(1):385–417, 2020.
- Kuhn, H. W. Extensive games. *Proceedings of the National Academy of Sciences*, 36(10):570–576, 1950.
- Kuhn, H. W. Extensive Games and the Problem of Information. *Annals of Mathematics Studies*, 28:193–216, 1953.
- Lanctot, M., Waugh, K., Zinkevich, M., and Bowling, M. Monte-Carlo sampling for regret minimization in extensive games. *Neural Information Processing Systems*, 2009.
- Lanctot, M., Lockhart, E., Lespiau, J.-B., Zambaldi, V., Upadhyay, S., Brodat, J., Srinivasan, S., Timbers, F., Tuyls, K., Omidshafiei, S., Hennes, D., Morrill, D., Muller, P., Ewalds, T., Faulkner, R., Krahan, J., De Vylder, B., Saeta, B., Bradbury, J., Ding, D., Borgeaud, S., Lai, M., Schrittwieser, J., Anthony, T., Hughes, E., Danihelka, I., and Ryan-Davis, J. OpenSpiel: A framework for reinforcement learning in games, 2019. URL <https://arxiv.org/abs/1908.09453>.
- Lattimore, T. and Szepesvári, C. *Bandit Algorithms*. Cambridge University Press, 2020. doi: 10.1017/9781108571401.
- Lee, C.-W., Kroer, C., and Luo, H. Last-iterate convergence in extensive-form games. *Neural Information Processing Systems*, 2021. URL <https://proceedings.neurips.cc/paper/2021/file/77bb14f6132ea06dea456584b7d5581e-Paper.pdf>.
- Littman, M. L. Markov games as a framework for multi-agent reinforcement learning. *International Conference on Machine Learning*, 1994.
- Liu, Q., Yu, T., Bai, Y., and Jin, C. A sharp analysis of model-based reinforcement learning with self-play. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 7001–7010. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/liu21z.html>.
- McAleer, S., Farina, G., Lanctot, M., and Sandholm, T. ESCHER: Eschewing importance sampling in games by computing a history value function to estimate regret. *CoRR abs/2206.04122*, 2022. doi: 10.48550/arXiv.2206.04122. URL <https://doi.org/10.48550/arXiv.2206.04122>.
- Moravčík, M., Schmid, M., Burch, N., Li, V., Morrill, D., Bard, N., Davis, T., Waugh, K., Johanson, M., and Bowling, M. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337):508–513, 2017. doi: 10.1126/science.aam6960. URL <https://www.science.org/doi/abs/10.1126/science.aam6960>.
- Munos, R., Brodat, J., Lespiau, J.-B., Rowland, M., De Vylder, B., Lanctot, M., Timbers, F., Hennes, D., Omidshafiei, S., Gruslys, A., Azar, M. G., Lockhart, E., and Tuyls, K. Fast computation of nash equilibria in imperfect information games. *International Conference on Machine Learning*, 2020.
- Nemirovski, A. Prox-Method with Rate of Convergence $O(1/t)$ for Variational Inequalities with Lipschitz Continuous Monotone Operators and Smooth Convex-Concave Saddle Point Problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- Nesterov, Y. Smooth minimization of non-smooth functions. *Mathematical programming*, 61(3):127–152, 2005.
- Neu, G. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. *Neural Information Processing Systems*, 2015.

- Pérolat, J., de Vylder, B., Hennes, D., Tarassov, E., Strub, F., de Boer, V., Muller, P., Connor, J. T., Burch, N., Anthony, T., McAleer, S., Elie, R., Cen, S. H., Wang, Z., Gruslys, A., Malysheva, A., Khan, M., Ozair, S., Timbers, F., Pohlen, T., Eccles, T., Rowland, M., Lanctot, M., Lespiau, J.-B., Piot, B., Omidshafiei, S., Lockhart, E., Sifre, L., Beauguerlange, N., Munos, R., Silver, D., Singh, S., Hassabis, D., and Tuyls, K. Mastering the game of Go with self-play. *Science* 2022. doi: 10.48550/ARXIV.2206.15378. URL <https://arxiv.org/abs/2206.15378>
- Romanovsky, J. V. Reduction of a game with complete memory to a matricial game. *Dokl. Akad. Nauk SSSR* 144:62–64, 1962.
- Schmid, M., Burch, N., Lanctot, M., Moravčík, M., Kadlec, R., and Bowling, M. Variance reduction in monte carlo counterfactual regret minimization (VR-MCCFR) for extensive form games using baselines. CoRR abs/1809.03057, 2018. URL <http://arxiv.org/abs/1809.03057>.
- Schmid, M., Moravčík, M., Burch, N., Kadlec, R., Davidson, J., Waugh, K., Bard, N., Timbers, F., Lanctot, M., Holland, Z., Davoodi, E., Christianson, A., and Bowling, M. Player of games. arXiv preprint arXiv:2112.03178 2021.
- Sidford, A., Wang, M., Yang, L., and Ye, Y. Solving discounted stochastic two-player games with near-optimal time and sample complexity. In Chiappa, S. and Candrandra, R. (eds.) *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020*, 26–28 August 2020, Online [Palermo, Sicily, Italy] volume 108 of *Proceedings of Machine Learning Research* pp. 2992–3002. PMLR, 2020. URL <http://proceedings.mlr.press/v108/sidford20a.html>
- Southey, F., Bowling, M., Larson, B., Piccione, C., Burch, N., Billings, D., and Rayner, C. Bayes' bluff: Opponent modelling in poker. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI)* pp. 550–558, 2005.
- Strens, M. A Bayesian Framework for Reinforcement Learning. In *International Conference on Machine Learning* 2000.
- Tammelin, O. Solving large imperfect information games using CFR+. arXiv preprint arXiv:1407.5042 2014.
- Tsallis, C. Possible Generalization of Boltzmann-Gibbs Statistics. *J. Statist. Phys.* 52:479–487, 1988. doi: 10.1007/BF01016429.
- von Neumann, J. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen* 100:295–320, 1928. ISSN 0025-5831; 1432-1807/e.
- von Stengel, B. Efficient computation of behavior strategies. *Games and Economic Behavior* 14(2):220–246, 1996.
- Waugh, K. and Bagnell, J. A. A unified view of large-scale zero-sum equilibrium computation. CoRR abs/1411.5007, 2014. URL <http://arxiv.org/abs/1411.5007>.
- Wei, C., Lee, C., Zhang, M., and Luo, H. Last-iterate convergence of decentralized optimistic gradient descent/ascent in finite-horizon competitive markov games. In Belkin, M. and Kpotufe, S. (eds.) *Conference on Learning Theory, COLT 2021*, 15–19 August 2021, Boulder, Colorado, USA volume 134 of *Proceedings of Machine Learning Research* pp. 4259–4299. PMLR, 2021. URL <http://proceedings.mlr.press/v134/wei21a.html>
- Wei, C.-Y., Hong, Y.-T., and Lu, C.-J. Online reinforcement learning in stochastic games. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.) *Advances in Neural Information Processing Systems* volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/36e729ec173b94133d8fa552e4029f8b-Paper.pdf>
- Xie, Q., Chen, Y., Wang, Z., and Yang, Z. Learning zero-sum simultaneous-move markov games using function approximation and correlated equilibrium. In Abernethy, J. and Agarwal, S. (eds.) *Proceedings of Thirty Third Conference on Learning Theory* volume 125 of *Proceedings of Machine Learning Research* pp. 3674–3682. PMLR, 09–12 Jul 2020. URL <https://proceedings.mlr.press/v125/xie20a.html>
- Zhang, B. H. and Sandholm, T. Finding and Certifying (Near-) Optimal Strategies in Black-Box Extensive-Form Games. In *AAAI Conference on Artificial Intelligence* 2021.
- Zhang, K., Kakade, S., Basar, T., and Yang, L. Model-based multi-agent rl in zero-sum markov games with near-optimal sample complexity. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.) *Advances in Neural Information Processing Systems* volume 33, pp. 1166–1178. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/0cc6ee01c82fc49c28706e0918f57e2d-Paper.pdf>
- Zhou, Y., Li, J., and Zhu, J. Posterior sampling for multi-agent reinforcement learning: solving extensive games with imperfect information. In *International Conference on Learning Representations* 2020. URL <https://openreview.net/pdf?id=Syg-ET4FPS>

Zimmert, J. and Seldin, Y. An optimal algorithm for stochastic and adversarial bandits. International Conference on Artificial Intelligence and Statistics 2019.

Zinkevich, M., Johanson, M., Bowling, M., and Piccione, C. Regret minimization in games with incomplete information. Neural Information Processing Systems 2007.

Appendix

Table of Contents

A Notation	14
B Extended related work	14
C Lower bound proof details	16
C.1 Variable action space sizes	16
C.2 Fixed action space size	20
D General tools	22
E Proof of Balanced FTRL	23
F Efficient updates and proof of correctness	33
G Proof of Adaptive FTRL	36

A. Notation

Notation	Meaning
S	state space of size $ S $
P_h	state-transition probability kernel at step h
r_h	reward function at step h
H	length of one episode
X	information set of the max-player of size $ X $
Y	information set of the min-player of size $ Y $
$A(x)$	action space at information set x of size $ A(x) $
$B(y)$	action space at information set y of size $ B(y) $
$A(X) = \bigcup_h A(X_h)$	union of information set-action space $A(X_h)$
$B(Y) = \bigcup_h B(Y_h)$	union of information set-action space $B(Y_h)$
$X_{h+1}(x_h; a_h)$	information sets of depth $h + 1$ that follow $(x_h; a_h)$
$\pi_h(j x_h)$	policy of the max-player at information-set x_h and step h
$\rho_h(j y_h)$	policy of the min-player at information-set y_h and step h
$\pi_{1:h}(x_h; a_h) = \prod_{h^0=1}^h \pi_{h^0}(a_{h^0} x_{h^0})$	realization plan of the max-player
$P_h(x_{h+1} x_h; a_h)$	probability to transition to x_{h+1} from $(x_h; a_h)$ with opponent
$r_h(x_h; a_h)$	average reward $r_h(x_h; a_h)$ when the opponent is
$p_{1:h}(x_h) = p_0(x_1) \prod_{h^0=1}^{h-1} P_{h^0}(x_{h^0+1} x_{h^0}; a_{h^0})$	adversarial realization plan against
$\ell_h(x_h; a_h) = p_{1:h}(x_h) (1 - r_h(x_h; a_h))$	adversarial loss against
T	number of episodes
R_{\max}^T	regret of max-player
$\ell^t = \sum_{k=1}^t \ell^k$	loss of the max-player at episode t
η	IX parameter
$\hat{\ell}_h^t = \sum_{k=1}^t \ell^k$	IX estimate of the loss of the max-player
$\hat{\ell}^t = \sum_{k=1}^t \hat{\ell}^k$	cumulative loss of the max-player
$p_{1:h}^t = \prod_{h^0=1}^h p_{1:h^0}(x_{h^0}; a_{h^0})$	induced probability distribution over $A(X_h)$ at depth h
$P_h^?$	balanced transition at step h
$\rho_h^?(x_h) = \prod_{h^0=1}^h \rho_{1:h^0}^?(x_{h^0})$	regularizers on $A(X_h)$
$D^t(\cdot \cdot)$	weighted (by η^t) dilated Shannon divergence D^t
$\eta^t(x_h; a_h)$	adaptive IX parameter $\eta^t(x_h; a_h)$
$\hat{P}_h^t(x_h; a_h)$	estimated transition $\hat{P}_h^t(x_h; a_h)$
$\hat{P}^t = \prod_{k=1}^t \hat{P}^k$	cumulative estimated transitions
$\alpha^t(x_h)$	adaptive learning rate α^t

Table 2. Table of notation use throughout the paper

B. Extended related work

In this section we review previous work on learning optimal strategies in IIGs.

Full feedback When the game is known, that is the information set structure space, transitions probability and reward function are provided a first line of work recasts the setting through the sequence-form representation of a game as a linear program which can be solved efficiently (Romanovsky, 1962; von Stengel, 1996; Koller et al., 1996). A second line of work relies on first-order optimization method for saddle point computation (Hoda et al., 2010; Kroer et al., 2015; 2018; 2020; Munos et al., 2020; Lee et al., 2021). In particular Hoda et al. (2010); Kroer et al. (2018) relies on the Nesterov smoothing technique Nesterov (2005) whereas Kroer et al. (2015; 2020) use the Prox algorithm (Nemirovski, 2004). These methods has a rate of convergence of $O(\frac{1}{\sqrt{t}})$ (poly($H; |A_X|; |B_Y|$)). Note that the rate is smaller than the lower bound of Theorem 3.1 since they assume full knowledge of the game. Even game dependent exponential rate can be obtained with these type of approach, see Gilpin et al. (2012) and Munos et al. (2020).

A third approach, counterfactual regret minimization (Zinkevich et al., 2007), leverages local regret minimisation, i.e. minimising a type of regret at each information set. Popular algorithms are based on the regret-matching algorithm (Hart & Mas-Colell, 2000; Gordon, 2007) such as CFR (Zinkevich et al., 2007) or based on a close variant of regret-matching, e.g. CFR+ (Tammelin, 2014; Burch et al., 2019; Farina et al., 2021a). Note that other local regret minimizers could be used, see for example Waugh & Bagnell (2014); Farina et al. (2019). These algorithms enjoy a guarantee of convergence of order $\mathcal{O}(\text{poly}(H; A_X; B_Y)) = \Theta(n^2)$.

Nevertheless, all the methods described above need the information set tree (or all the state space) in order to compute one update. The cost of one traversal is of order $\mathcal{O}(X + Y)$ if the transitions and the actions of the other player are sampled; see for example the external-sampling CFR algorithm (Lanctot et al., 2009).

Trajectory feedback A way to tackle the aforementioned issues is to consider the agnostic setting where the agent has no prior knowledge of the game and only observes trajectories of the game, i.e. precisely, the rewards, the transition probabilities are unknown. This is the setting considered in this paper.

Model-based A first method to deal with this limited feedback is to build a model of the game then run any full feedback algorithm in this model. For example, Zhou et al. (2020) use posterior sampling (PS, Strens, 2000) to learn a model and then use the CFR algorithm in games sampled from the posterior. They obtain a convergence rate of order $\mathcal{O}(\text{poly}(H; S; A; B)) = \Theta(n^2)$ but only when the games are actually sampled according to the known prior. Instead, Zhang & Sandholm (2021) rely on the principle of optimism in presence of uncertainty to incrementally build a model of the game. The CFR algorithm is fed with optimistic estimates of the local regrets. They prove a high-probability sample complexity of order $\mathcal{O}(\text{poly}(H; S; A; B)) = \Theta(n^2)$.

Model-free Another line of work (Lanctot et al., 2009; Johanson et al., 2012; Schmid et al., 2018; Farina et al., 2020) directly estimate the local regret via importance sampling that are then fed to the CFR algorithm. In particular, the outcome-sampling MCFR (Lanctot et al., 2009; Farina et al., 2020) builds an importance sampling estimate of the counterfactual regret by playing according to a well-chosen balanced policy. Intuitively, this policy should ensure to explore all the information sets. Note that, depending on the structure of the information set space, playing uniformly over the actions at each information set is not necessarily a good choice. Instead, Farina et al. (2020) propose as a balanced policy to play action with probability proportional to the number of leaves in the sub-tree of possible next information sets. In particular, the outcome-sampling MCFR algorithm requires the knowledge of the information set space structure to build its balanced policy. Nonetheless, in order to obtain optimal strategies with high probability, MCFR needs at most $\mathcal{O}(H^3(A_X + B_Y)) = \Theta(n^3)$ realizations of the game (Farina et al., 2020; Bai et al., 2022).

Latter, Kozuno et al. (2021) propose to combine Online Mirror Descent (OMD) with dilated Shannon entropy as regularizer and importance sampling estimate of the losses of a player, see also Farina et al. (2021b). They prove a sample complexity, for the proposed algorithm, of order $\mathcal{O}(H^2(XA_X + YB_Y)) = \Theta(n^2)$. Interestingly, they do not need to know in advance the structure of the information set space to obtain this bound. However, the sample complexity of OMD does not match the lower bound for this setting which is of order $\mathcal{O}((A_X + B_Y)) = \Theta(n)$, see Section 3 and Bai et al. (2022). Recently, Bai et al. (2022) propose the Balanced OMD algorithm that also relies on OMD but with a dilated entropy weighted by the realization plans of balanced policies as regularizers. These balanced policies generalize the one introduced by Farina et al. (2020) for all depths in the information set tree. Balanced OMD they prove a sample complexity of order $\mathcal{O}(H^3(A_X + B_Y)) = \Theta(n^3)$ with an improved dependency on the sizes of the information set spaces at the price of a worse dependence in the horizon.

Perfect information Markov game Another line of work consider Markov game Kuhn (1953) with perfect information and limited feedback. However, they do not assume perfect recall. Sidford et al. (2020); Zhang et al. (2020); Daskalakis et al. (2020); Wei et al. (2021) consider the case where generative models are available whereas Wei et al. (2017); Bai et al. (2020); Xie et al. (2020); Liu et al. (2021) deal with the trajectory feedback case. Although this setting is related to ours there is no direct comparison between the two.

⁶They only need to know the size of the information set spaces to tune optimally the learning rate

C. Lower bound proof details

To formally state lower bounds, we start with a tad of additional notation. For the proofs of the lower bounds, we consider games in which one player (say the min-player) has no effect on both rewards and state-transition dynamics. Therefore, we omit random variables related to the min-player.

An algorithm is a sequence $\pi := (\pi^t)_{t=1}^{T+1}$, where π^t is a measurable mapping from $(X)^{(t-1)H} \times R^{(t-1)H}$ to \mathcal{A}_{\max} . In t -th episode, the algorithm outputs a policy $\pi^t = \pi^t(H_t)$ given the history $H_t := (x_h^u; a_h^u; r_h^u)_{u=1; h=1}^{t-1; H}$ and follows it. By Ionescu–Tulcea theorem (Lattimore & Szepesvári, 2020, Theorem 3.3), there exists a probability space consistent with this procedure. After T -th episode, the algorithm outputs a final policy $\pi^{T+1} = \pi^{T+1}(H_{T+1})$.

We prove the following theorems. Note that we assume deterministic rewards, but lower bounds show that the difficulty of IIGs is the same as the difficulty of bandit problems when $\epsilon = 1$. This is because we made use of the partial observability such that rewards look stochastic from the view point of players.

Theorem C.1 (Regret Lower Bound; Variable Action Space Size) Fix horizon H , number of episodes E , total number of actions A_X , and a positive scalar $\epsilon \in (0; 1)$. If $A_X < H$, there is no such game. If $A_X \leq H$, $\epsilon < 1/4$, and $T \geq 0.4K \log(4/\epsilon)$, for any algorithm, there exists a game such that with probability greater than $1/4$, the algorithm suffers from the regret of

$$\frac{1}{4} \frac{(A_X - H) \min_{A_X \leq H} \log(1/\epsilon)}{H} :$$

Theorem C.2 (Sample Complexity Lower Bound; Variable Action Space Size) Fix horizon H , total numbers of actions $A_X; B_Y$, and positive scalars $\epsilon \in (0; 1)$; $\epsilon' \in (0; H)$. If either $A_X < H$ or $B_Y < H$, there is no such game. If $A_X \leq H$, $B_Y \leq H$, and $\epsilon < 1/4$, for any algorithm, there exists a game such that the algorithm needs at least

$$\frac{(A_X - H) \min_{A_X \leq H} \log(1/\epsilon) + (B_Y - H) \min_{B_Y \leq H} \log(1/\epsilon')}{\epsilon'^2} \log(1/\epsilon)$$

episodes to output ϵ -NE with probability at least $1/4$.

Theorem C.3 (Regret Lower Bound; Fixed Action Space Size) Fix horizon H , number of episodes E , number of actions A at each information set, number of information sets X, Y , and a positive scalar $\epsilon \in (0; 1)$. If $X < (A^H - 1)/(A - 1)$, there is no such game. If $X \leq (A^H - 1)/(A - 1)$, $\epsilon < 1/4$, and $T \geq 0.4K \log(4/\epsilon)$, for any algorithm, there exists a game such that with probability greater than $1/4$, the algorithm suffers from the regret of

$$\frac{1}{4} \frac{X \log(1/\epsilon)}{T A_X} :$$

Theorem C.4 (Sample Complexity Lower Bound; Fixed Action Space Size) Fix horizon H , number of actions $A; B$ at each information set, numbers of information sets X, Y , and positive scalars $\epsilon \in (0; 1)$; $\epsilon' \in (0; H)$. If either $X < (A^H - 1)/(A - 1)$ or $B < (B^H - 1)/(B - 1)$, there is no such game. If $X \leq (A^H - 1)/(A - 1)$, $(B^H - 1)/(B - 1) \leq Y < (B^H - 1)/(B - 1)$, and $\epsilon < 1/4$, for any algorithm, there exists a game such that the algorithm needs at least

$$\frac{A_X + B_Y}{\epsilon'^2} \log(1/\epsilon)$$

episodes to output ϵ -NE with probability at least $1/4$.

C.1. Variable action space sizes

We first consider a case in which the number of viable actions depends on information sets.

Proof sketch In the hard game instance we consider, we make actions of one player ineffective depending on which one of A_X and B_Y is bigger. Specifically, if $A_X > B_Y$, the min-player's actions have no effect on the state-transition dynamics and rewards.⁷ Then the game reduces to a Markov decision process, and finding an optimal policy is equivalent to finding an ϵ -NE. For the time being, we assume that $A_X > B_Y$.

⁷In case of ties, make actions of the min-player ineffective.

Due to the perfect-recall assumption (at least one successive information set for each action except the final step), A_x must be greater than or equal to H . Furthermore, when $A_x = H$, there is only one policy, and the sample complexity lower bound is trivially $0 = A_x - H$. Hereafter we assume that $A_x > H$.

Figure 2. A hard game instance with information-set-dependent action space from the view point of the max-player.

Figure 2 depicts hard game instances under different conditions. White circles are states, and states at the top (and identified by integers) are initial states. Edges from each state are viable actions at a corresponding state. Blue circles are rewarding actions, where the max-player receives the reward 1 . Red circles are non-rewarding actions. Shaded boxes are information sets, and all states in a box belong to the same information set.

When $A_x = 2H$, we construct a game shown in the bottom left panel. For brevity let $b_{A_x} = H$ and assume that A_x is divisible by H . In this case there is no right-most branch. If A_x is not divisible by H , append a small branch as shown in the figure and set state-transition probability to the small branch to 0 . In this game, there are 2^k initial states, at which there are A_x viable actions. Each initial state is identified by an integer from 0 to $2^k - 1$. An initial state is sampled as follows: in the beginning of each episode, k binary integers are sampled from (ordered) Bernoulli distributions and concatenated to construct a binary representation of an integer from 0 to $2^k - 1$ (e.g., 010111), which is an initial state in this episode. The parameter of Bernoulli distribution is denoted by p_k . The reward value for k -th action is the k -th bit of the initial state in binary representation. After the first step, the state transitions in a deterministic way as shown in the figure, and reward values at later time steps are the same as the reward value at the first time step.

To get an intuition about this hard game instance, it is beneficial to consider a case where $A_x = H$. In this case, the game essentially reduces to a bandit problem shown in the top left panel of Figure 2. In this case the regret lower bound of (\sqrt{KT}) and the sample complexity lower bound ($k = \sqrt{2}$) are known to hold. When the reward is scaled by the lower bounds become $(\sqrt{H(A_x - H)T})$ and $(\sqrt{H(A_x - H)k})$, respectively. The hard game instance described above instantiates this idea.

The hard game construction described above does not work when $2H > A_x > H$. In this case, we construct a

game shown in the bottom right panel. The idea is almost the same, but there are only $2H - A_X + 1$ states, and the player makes decision at $(2H - A_X + 1)$ -th time step instead of the first time step. There the max-player has 2 actions, and reward values at $(2H - A_X + 1)$ -th time step and thereafter are determined as before. This game is essentially a Bernoulli bandit with reward scaled by $A_X - H$. In this case the regret lower bound $(A_X - H) \sqrt{2T}$ and the sample complexity lower bound of $(2(A_X - H))^2$ are known to hold.

If $A_X > B_Y$ does not hold, we can derive the same result with A_X replaced by B_Y by exchanging the role of max- and min-players in the previous argument. Consequently, depending on which one of A_X or B_Y is bigger, we can use A_X or B_Y in lower bounds. The claimed results follow from a fact that $\max\{a, b\} = \frac{a+b}{2} + \frac{|a-b|}{2}$.

Full Proof. We assume that $A_X \geq 2H$ as other cases can be handled similarly. We consider $Hc + 1$ games and identify each game by an integer from 0 to $bA_X = Hc$. In 0-th game, we set α_k to 0.5. In k -th game, we set α_k to 0.5, and p_k to $0.5 + \frac{r}{2v}$, where $r \in (0, 0.5)$ is specified later in the proof. In other words, an optimal arm in k -th game is k -th action. We let \mathbb{P}_k denote the probability measure induced by the interconnection of the algorithm and game. Expectation under \mathbb{P}_k is denoted by E_k . We also let \mathbb{P}_k^0 and E_k^0 denote the law of $(x_h^u; a_h^u; r_h^u)_{u=1; h=1}^{T; H}$ and expectation under \mathbb{P}_k^0 when the algorithm is run in game k .

The proofs rely on the following core lemma.

Lemma C.5. Fix scalars $r \in (0, 1)$ and $v \in (0, 1)$. For brevity let $K := bA_X = Hc$. If $\alpha \in (0, 3)$, for any measurable function $f : A(X)^{TH} \rightarrow \mathbb{R}$, $1g^{TH} \rightarrow \mathbb{R}$ and $k \in [K]$, it holds that

$$\mathbb{P}_k(f(H_{T+1}) = 1) > \exp(-s - v) \mathbb{P}_0(f(H_{T+1}) = 1) - \frac{4}{v} \alpha^2 n_k;$$

where $s := \frac{4}{3} \log \frac{1}{1-r} + \frac{r}{2v} \log \frac{1}{1-2r}$ and $n_k := E_0 \prod_{t=1}^T I_{f_{a_t^i=k}}$.

Proof. Let E be an event $f(H_{T+1}) = 1$. We begin with rewriting $\mathbb{P}_k(E)$ as follows:

$$\mathbb{P}_k(E) = E_k[f(H_{T+1})] = E_k^0[f(H_{T+1})] = E_0^0[f(H_{T+1})] \frac{d\mathbb{P}_k^0}{d\mathbb{P}_0^0}(H_{T+1}); \tag{1}$$

We derive a lower bound for $\frac{d\mathbb{P}_k^0}{d\mathbb{P}_0^0}(H_{T+1})$. Because $A(X)^{TH} \rightarrow \mathbb{R}$, $1g^{TH}$ is discretized, $\frac{d\mathbb{P}_k^0}{d\mathbb{P}_0^0}(H_{T+1})$ is simply an importance sampling ratio. Furthermore, changing the game k only changes the reward distribution when k -th action is taken. Therefore,

$$\begin{aligned} \frac{d\mathbb{P}_k^0}{d\mathbb{P}_0^0}(H_{T+1}) &= \prod_{t=1}^T r_1^t \frac{0.5 + I_{f_{a_t^i=k}}}{0.5} + (1 - r_1^t) \frac{0.5 - I_{f_{a_t^i=k}}}{0.5} \\ &= \prod_{t=1}^T r_1^t (1 + 2I_{f_{a_t^i=k}}) + (1 - r_1^t) (1 - 2I_{f_{a_t^i=k}}) = \prod_{t=1}^T (1 + P_t); \end{aligned}$$

where $P_t := 2(2r_1^t - 1)I_{f_{a_t^i=k}}$. From a fact that $1 + x \geq \exp(x) \geq x^2$ for $x \in [0, 0.6]$, we deduce that

$$\frac{d\mathbb{P}_k^0}{d\mathbb{P}_0^0}(H_{T+1}) \geq \exp(-S_T - V_T);$$

where we define $S_T := \prod_{t=1}^T P_t$ and $V_T := \prod_{t=1}^T P_t^2 = 4^{-T} \prod_{t=1}^T I_{f_{a_t^i=k}}$.

Returning to equation (1), it clearly holds that

$$\begin{aligned} \mathbb{P}_k(E) &= E_0^0[f(H_{T+1})] I_{f_{S_T < s; V_T < v}} \frac{d\mathbb{P}_k^0}{d\mathbb{P}_0^0}(H_{T+1}) \\ &> \exp(-s - v) E_0^0[f(H_{T+1})] I_{f_{S_T < s; V_T < v}} \\ &= \exp(-s - v) E_0[f(H_{T+1})] I_{f_{S_T < s; V_T < v}}; \end{aligned}$$

We lower-bound $f(H_{T+1})|_{f_{S_T < s; V_T > v}}$. To this end, note that $A \setminus B = (A \cap B^c) \cup (A \setminus B)$ for any sets A and B . Therefore,

$$f(H_{T+1})|_{f_{S_T < s; V_T > v}} = f(H_{T+1}) \cdot \mathbb{1}_{f_{S_T < s; V_T > v}} = f(H_{T+1}) \cdot \mathbb{1}_{f_{V_T > v}} + f(H_{T+1}) \cdot \mathbb{1}_{f_{S_T < s; V_T > v}} \cdot \mathbb{1}_{f_{V_T \leq v}}$$

and thus,

$$E_0 f(H_{T+1})|_{f_{S_T < s; V_T > v}} \geq P_0(E) \cdot \underbrace{P_0(S_T < s; V_T > v)}_{\text{by (a)}} + \underbrace{P_0(V_T > v)}_{E_0[V_T]=v \text{ by (b)}}$$

It is easy to see that (b) follows from Markov's inequality. To see why (a) follows, let \mathcal{F}_t be the σ -algebra generated by $(s_h^u; a_h^u)_{u=1}^t; h=1$ and a_1^t . Then, (a) follows from Freedman's inequality by noting that $(F_t)_{t=1}^T$ is a martingale difference sequence with respect to the filtration $(\mathcal{F}_t)_{t=1}^T$, F_t is almost surely bounded by v , and $E_0[F_t | \mathcal{F}_{t-1}] = 0$. Consequently,

$$P_k(E) > \exp(-\frac{4v^2}{v}) P_0(E) = \frac{4}{v} E_0 \sum_{t=1}^T \mathbb{1}_{f_{a_t^1 = k}} = \frac{4}{v} n_k$$

This concludes the proof. □

Regret Lower Bound. Now we are ready to prove the regret lower bound. For each $k \in [K]$, let E_k be an event $\prod_{t=1}^T \mathbb{1}_{f_{(a_k^j)x_1^t} = 0; 5T}$. Note that while implicit, $\prod_{t=1}^T \mathbb{1}_{f_{(a_k^j)x_1^t} = 0}$ is a function of H_{T+1} , and thus E_k can be rewritten in a form of $f(H_{T+1}) = \mathbb{1}_{E_k}$. As we verify later, our choice of v is smaller than 0.3 . Therefore from the previous lemma,

$$\max_{k \in [K]} P_k(E_k) \geq \frac{1}{K} \sum_{k \in [K]} P_k(E_k) > \exp(-\frac{4v^2}{v}) \frac{1}{K} \sum_{k=1}^K P_0(E_k) \geq \frac{4T^2}{Kv}$$

Note that $\sum_{k \in [K]} P_0(E_k) = E_0 \sum_{k \in [K]} \mathbb{1}_{f_{E_k} = 1} = K \cdot \mathbb{1}_{f_{E_k} = 0 \text{ for some } k} = 1$ since $f_{E_k} = 0$ for some k implies that $f_{E_j} = 1$ for $j \neq k$. Thus,

$$\max_{k \in [K]} P_k(E_k) > \exp(-\frac{4v^2}{v}) \frac{1}{K} \geq \frac{4T^2}{Kv}$$

Setting $v = \frac{32T^2}{K}$ and $\frac{1}{v} = \frac{1}{8}$,

$$\max_{k \in [K]} P_k(E_k) > \frac{1}{4} \exp(-4 \log 2) \geq \frac{1}{8} \frac{3T^2 \log 2}{K} \geq \frac{32T^2}{K}$$

where we used an assumption that $K \geq 2$. Equating the right hand side to $\frac{1}{8}$ and solving for T , we choose

$$T \geq \sqrt{\frac{K}{16T} \log 2 + \frac{1}{8} \frac{3K}{T} \log 2} + \frac{K}{32T} \log \frac{1}{4} \geq \frac{K}{16T} \log 2 + \frac{1}{8} \frac{3K}{T} \log 2 < \frac{K}{32T} \log \frac{1}{4} < 0.3; \quad (2)$$

where the middle inequality follows since $a + b < \sqrt{a^2 + b^2}$ for $a, b \geq 0$. Then it holds that $\max_{k \in [K]} P_k(E_k) > \frac{1}{8}$. Because the event E_k implies that the regret is greater than or equal to $\frac{1}{8} H$ in k -th game, we have a lower bound $(\frac{1}{8} H^2 T K \log(4)) = (\frac{1}{8} H^2 T (A_X = H) \log(4)) = (\frac{1}{8} H T (A_X = H) \log(4))$.

Sample Complexity Lower Bound. Now we are ready to prove the sample complexity lower bound. Let E_k be an event $\prod_{t=1}^{T+1} \mathbb{1}_{f_{(a_k^j)x_1^t} = 0; 5}$ for each $k \in [K]$. By the same argument as the one in the regret lower bound proof, $\max_{k \in [K]} P_k(E_k) > \frac{1}{8}$ for T in Equation 2. Since in k -th game, the event E_k implies that the simple regret of π^{T+1} is greater than or equal to $\frac{1}{8} H$, we have a lower bound for the simple regret of π^{T+1} $(\frac{1}{8} H (A_X = H) = T \log(4))$. Therefore unless $T \geq \frac{1}{8} H (A_X = H) = 2 \log(4)$, the algorithm fails to output-optimal policy with probability greater than $\frac{1}{8}$.

C.2. Fixed action space size

Now we turn to the case where the action space depends on information sets. Since the proof is relatively easy, we directly give a full proof.

In the hard game instance we will consider, we make actions of one player ineffective depending on which A and B_γ is bigger, as before. For the time being, we assume that B_γ . Furthermore we restrict rewards to be binary but allow them to be stochastic. As we did in the proof of Lemma C.5, we can create a game in which rewards are deterministic but look stochastic from the view point of players.

First of all, we claim that there is no game such that $(A^{H-1})=(A-1)$. This is because $x-1+A+x+A^{H-1}=(A^{H-1})=(A-1)$ due to the perfect-recall assumption (at least one successive information set for each action except the final time step). Hereafter we assume that $(A^{H-1})=(A-1)$.

Figure 3. A hard game instance with information-set-independent action space from the view point of the max-player.

Figure 3 depicts the hard game instance under different conditions. In all cases, all information sets are singleton. Unless a number is given, all edges from action nodes to state nodes mean deterministic transitions. Numbers along edges indicate state-transition probabilities of the edges. In all cases, the max-player receives reward at action nodes with some probability specified later in the proof.

The top left panel shows the case where $(A^{H-1})=(A-1)$. In this case, the game tree is a binary tree with deterministic transitions. This case is dealt by Bai et al. (2022), and we obtain a similar (pseudo-)regret bound.

The top right panel shows the case where $(A^{H-1})=(A-1)+1$. In this case, an additional A -ary tree with depth 1 is appended at the left-most action node at $(H-1)$ -th time step. State-transition probabilities to the successor state nodes are 0.5. Up to $(A^{H-1})=(A-1)+A^{H-1}$, additional A -ary trees with depth 1 are added similarly to different action nodes at $(H-1)$ -th time step one by one from left to right action nodes. For example, if $(A^{H-1})=(A-1)+2$, another additional A -ary tree with depth 1 is appended at the second left-most action node at $(H-1)$ -th time step.

The bottom panel shows the case where $(A^{H-1})=(A-1)+A^{H-1}+1$. In this case, we remove the additional A -ary tree at the left-most action node at $(H-1)$ -th time step and append an A -ary tree with depth 2 at the left-most

action node at $(H-2)$ -th time step. Up to $X = (A^{H-1}) = (A-1) + A^{H-1} + A^{H-2}$, additional A -ary trees with depth 1 are added similarly to different action nodes at $(H-2)$ -th time step one by one from left to right action nodes. This procedure can be repeated until all action nodes at the first time step have additional branches. At this point, there are $(A^{H-1}) = (A-1) + A^{H-1} + \dots + A = 2(A^{H-1}) = (A-1) - 1$ information sets.

Now we specify reward probabilities at leaf action nodes and construct a $\min_X (A^{H-1}) = (A-1); A^{H-1}g + 1$ game. We index leaf action nodes by integers from $2A^H$. We let a_i and x_i denote the corresponding action at node i and its predecessor information set (i.e., an information set at which can be taken). In k -th game, we set all reward probabilities to 0.5 . In k -th game, we set reward probabilities of all leaf action nodes except one to 0.5 , and the reward probability of k -th leaf action node to $0.5 + \epsilon$, where $\epsilon \in (0; 0.5)$ is specified later in the proof. In other words, an optimal leaf action node in k -th game is k -th one. We let P_k denote the probability measure induced by the interconnection of the algorithm and game k . Expectation under P_k is denoted by E_k . We also let P_k^0 and E_k^0 denote the law of $(x_h^u; a_h^u; r_h^u)_{u=1;h=1}^{T;H}$ and expectation under P_k^0 when the algorithm is run in game k .

By using the almost same proof as that of Lemma C.5, we can deduce the following lemma.

Lemma C.6. Fix scalars $\epsilon \in (0; 1)$ and $\nu \in (0; 1)$. For brevity let $K := 2A^H$. If $\epsilon > 0.3$, for any measurable function $f : A(X)^{T;H} \rightarrow [0; 1]$, $g : [0; 1] \rightarrow [0; 1]$ and $k \in [K]$, it holds that

$$P_k(f(H_{T+1}) = 1) > \exp(-\epsilon - \nu) P_0(f(H_{T+1}) = 1) - \frac{4\epsilon^2}{\nu} n_k;$$

where $n_k := \frac{4}{3} \log \frac{1}{\epsilon} + \frac{r}{2\nu \log 2} \sum_{t=1}^T \sum_{i \in I_{f_{x_H^t} = x_k, a_H^t = a_k} g}$.

Regret Lower Bound. Now we are ready to prove the regret lower bound. For each $k \in [K]$, let E_k be an event $\prod_{t=1}^T \prod_{h=1}^H (x_k; a_k) \in [0; 0.5T]$. Note that while implicit, $\prod_{t=1}^T \prod_{h=1}^H (x_k; a_k)$ is a function of H_{T+1} , and thus E_k can be rewritten in a form of $(H_{T+1}) = 1$. As we verify later, our choice of ϵ is smaller than 0.3 . Therefore from the previous lemma,

$$\max_{k \in [K]} P_k(E_k) - \frac{1}{K} \sum_{k \in [K]} P_k(E_k) > \exp(-\epsilon - \nu) \frac{1}{K} \sum_{k=1}^K P_0(E_k) - \frac{4T\epsilon^2}{K\nu};$$

Note that $\sum_{k \in [K]} P_0(E_k) = E_0[\sum_{k \in [K]} I_{f_{E_k} g}] \leq K A^{H-1}$. To see why it is the case, suppose that $I_{f_{E_k} g} = 0$ for some k , and let $(x_h; a_h)_{h=1}^H$ be information sets and actions from the root to $(x_k; a_k)$. Then $\prod_{t=1}^T \prod_{h=1}^H (x_k; a_k) > 0.5T$ means that $\prod_{t=1}^T \prod_{h=1}^H (x_h; a_h) > 0.5T$ for all h since $\prod_{t=1}^T \prod_{h=1}^H (x_h; a_h) = \prod_{t=1}^T \prod_{h=1}^H (x_{h-1}; a_{h-1}) \prod_{h=1}^H (a_h | x_h)$. Furthermore it also holds that $\prod_{t=1}^T \prod_{h=1}^H (a_h | x_h) > 0.5T$ for any h . If there is no branch along the path to the leaf action node to the root node, every trajectory to leaf action nodes certainly contains $(a_h | x_h)$. Therefore $\prod_{t=1}^T \prod_{h=1}^H (a_h | x_h) > 0.5T$ holds for any $a_1 \in A$ and $a_1 g$, and $I_{f_{E_k} g} = 0$ for some k implies that $I_{f_{E_j} g} = 1$ for some $j \in [K]$. If there is an additional branch, the same argument holds for all leaf action nodes except those in the branch. In the worst case, the number of leaf action nodes in a branch is A^{H-1} . Thus,

$$\max_{k \in [K]} P_k(E_k) > \exp(-\epsilon - \nu) \left(1 - \frac{1}{A}\right) \frac{4T\epsilon^2}{K\nu};$$

Setting $\nu = \frac{32T\epsilon^2}{K}$ and $\epsilon = \frac{1}{8}$,

$$\max_{k \in [K]} P_k(E_k) > \frac{1}{4} \exp(-4 \log 2 - 8) \frac{3T}{K} \log 2 - \frac{32T\epsilon^2}{K};$$

where we used an assumption that $\epsilon > 0.3$. Equating the right hand side to ϵ and solving for ϵ , we choose

$$\epsilon = \sqrt{\frac{\nu}{\frac{K}{16T} \log 2 + \frac{1}{8} \frac{3K}{T} \log 2}} + \frac{K}{32T} \log \frac{1}{4} - \frac{K}{16T} \log 2 - \frac{1}{8} \frac{3K}{T} \log 2 < \sqrt{\frac{K}{32T} \log \frac{1}{4}} < 0.3; \quad (3)$$

where the middle inequality follows since $\frac{p}{a+b} < \frac{p}{a} + \frac{p}{b}$ for $a, b \in (0, 1)$. Then it holds that $\max_{k \in [K]} P_k(E_k) > \frac{1}{2}$. Because the event E_k implies that the regret is greater than or equal to $\frac{1}{2} H$ in k -th game, we have a lower bound $(\frac{1}{2} T A^H \log(4)) = (\frac{1}{2} T X A \log(4))$, where $X = 2(1 + A + \dots + A^{H-1}) = X_H(1 + 1 = A + \dots + 1 = A^{H-1})$. $2X_H$ is used.

Sample Complexity Lower Bound. Now we are ready to prove the sample complexity lower bound. Let E_k be an event $\frac{1}{2} \mathbb{1}_{\{R_k \geq H\}}$ for each $k \in [K]$. By the same argument as the one in the regret lower bound proof, $\max_{k \in [K]} P_k(E_k) > \frac{1}{2}$ for H in Equation 3. Since in k -th game, the event E_k implies that the simple regret of π^{k+1} is greater than or equal to $\frac{1}{2} H$, we have a lower bound for the simple regret of $(\frac{1}{2} X A = T \log(4))$. Therefore unless $T \frac{1}{2} X A = \frac{1}{2} T \log(4)$, the algorithm fails to output ϵ -optimal policy with probability greater than

D. General tools

All proofs of the following sections will focus on the max player. We will denote $(X_t)^t$ its history at episode t and use $(F_t)_{t \in [0; T]}$ the filtration such that $F_0 = \mathcal{F}; g, F_t = \mathcal{F}_t; r^1; A(X)^1; \dots; r^t; A(X)^t; \dots; r^T; A(X)^T$ is the filtration of \mathcal{F} -algebra generated by the relevant random variables for the max-player up to the round $t+1$.

We first provide a slight generalization of Lemma 1 by Neu (2015) for our settings, as not fixed in advance in the adaptive case.

Lemma D.1. Let $h \in [H], (x_h; a_h) \in A(X_h), \tau_h(x_h; a_h)$ a stopping time with respect to $(\mathcal{F}_t)_{t \in [0; T]}$, and $\frac{1}{h}(x_h; a_h) > 0$ a fixed constant such that for all $(x_h; a_h), \tau_h(x_h; a_h) \leq \frac{1}{h}(x_h; a_h)$. Then, with probability $1 - \epsilon$,

$$E_h^{(x_h; a_h)}(X_h; a_h) \leq L_h^{(x_h; a_h)}(x_h; a_h) \frac{\log(1 + \frac{1}{h}(x_h; a_h))}{2 \frac{1}{h}(x_h; a_h)} :$$

Proof. We first define the random process $(S^t)_{t \in [0; T]}$ with respect to $(F_t)_{t \in [0; T]}$

$$S^t := \exp \left(\sum_{k=1}^t \frac{1}{h}(x_h; a_h) \left(r_h^k(x_h; a_h) - \frac{1}{h}(x_h; a_h) \right) \right) :$$

With the inequality $\frac{z}{1+z} \leq \log(1+z)$ for $z \geq 0$, we have for all $(x_h; a_h)$, with $\tau_{1:h}^{-t} := \tau_{1:h}(x_h; a_h)$,

$$\begin{aligned} 2 \frac{1}{h}(x_h; a_h) E_h^{(x_h; a_h)}(X_h; a_h) &= 2 \frac{1}{h}(x_h; a_h) \frac{1}{\tau_{1:h}^{-t} + \frac{1}{h}(x_h; a_h)} \mathbb{E} \left[\sum_{k=1}^{\tau_{1:h}^{-t}} r_h^k(x_h; a_h) \mid x_h = x_h^t; a_h = a_h^t \right] \\ &= 2 \frac{1}{h}(x_h; a_h) \frac{1}{\tau_{1:h}^{-t} + \frac{1}{h}(x_h; a_h)} \frac{1}{1 - r_h^t} \mathbb{E} \left[\sum_{k=1}^{\tau_{1:h}^{-t}} r_h^k(x_h; a_h) \mid x_h = x_h^t; a_h = a_h^t \right] \\ &= \frac{2 \frac{1}{h}(x_h; a_h) (1 - r_h^t)^{-\tau_{1:h}^{-t}}}{1 + \frac{1}{h}(x_h; a_h) (1 - r_h^t)^{-\tau_{1:h}^{-t}}} \mathbb{E} \left[\sum_{k=1}^{\tau_{1:h}^{-t}} r_h^k(x_h; a_h) \mid x_h = x_h^t; a_h = a_h^t \right] \\ &= \log(1 + 2 \frac{1}{h}(x_h; a_h) \frac{1}{h}(x_h; a_h)) : \end{aligned}$$

Which gives, using $\log(1+z) \leq e^z$,

$$\begin{aligned} E_h \exp \left(\sum_{k=1}^t \frac{1}{h}(x_h; a_h) \left(r_h^k(x_h; a_h) - \frac{1}{h}(x_h; a_h) \right) \right) &\leq E_h \left(1 + 2 \frac{1}{h}(x_h; a_h) \frac{1}{h}(x_h; a_h) \right)^{\tau_{1:h}^{-t}} \\ &= 1 + 2 \frac{1}{h}(x_h; a_h) \frac{1}{h}(x_h; a_h) \tau_{1:h}^{-t} \\ &\leq \exp \left(2 \frac{1}{h}(x_h; a_h) \frac{1}{h}(x_h; a_h) \tau_{1:h}^{-t} \right) : \end{aligned}$$

We thus have that $(S^t)_{t \in [0; T]}$ is a super-martingale, and using the Markov inequality we get

$$\mathbb{P} \left[\sum_{t=1}^T \left(u_t - \mathbb{E}(u_t | \mathcal{F}_{t-1}) \right) \geq H \right] \leq \exp \left(-\frac{2H^2 \log(1/\delta)}{4H^2 T} \right) = \delta.$$

□

We also state a simple consequence of Azuma-Hoeffding inequality that we will use multiple times in the following sections.

Lemma D.2. Let $(u_t)_{t \in [T]}$ be a random process adapted to $(\mathcal{F}_t)_{t \in [T]}$ such that $0 \leq u_t \leq H$ for all $t \in [T]$. Then, with a probability at least $1 - \delta$,

$$\sum_{t=1}^T \left(u_t - \mathbb{E}(u_t | \mathcal{F}_{t-1}) \right) \leq H \sqrt{2T \log(1/\delta)}.$$

The same property can also be shown for $\sum_{t=1}^T \left(\mathbb{E}(u_t | \mathcal{F}_{t-1}) - u_t \right)$.

Proof. We define the martingale $(M_t)_{t \in [0; T]}$ adapted to $(\mathcal{F}_t)_{t \in [0; T]}$ with

$$M_t := \sum_{k=1}^t \left(u_k - \mathbb{E}(u_k | \mathcal{F}_{k-1}) \right).$$

We have, thanks to the hypothesis, for all $t \in [T]$, that $-H \leq M_t - M_{t-1} \leq H$. Azuma-Hoeffding inequality then allows us to conclude

$$\mathbb{P} \left[M_T \geq H \sqrt{2T \log(1/\delta)} \right] \leq \exp \left(-\frac{2H^2 T \log(1/\delta)}{4H^2 T} \right) = \delta.$$

The proof for $\sum_{t=1}^T \left(\mathbb{E}(u_t | \mathcal{F}_{t-1}) - u_t \right)$ is exactly the same. □

Finally, we state a lemma useful for the analysis of **Adaptive FTRL**.

Lemma D.3. Let $(u_t)_{t \in [T]}$ be a non-negative sequence that verifies for all $t \in [T]$, $u_t \leq 1 + U_{t-1}$ where $U_t = \sum_{k=1}^t u_k$. Then

$$\sum_{t=1}^T \frac{u_t}{1 + U_{t-1}} \leq \log_2(1 + U_T).$$

Proof. Let $\tilde{u}_t = \frac{u_t}{1 + U_{t-1}}$. Using the concavity of \log_2 , we get $\tilde{u}_t \leq \log_2(1 + \tilde{u}_t)$ as $\tilde{u}_t \in [0; 1]$ by assumption. Summing over t then yields

$$\sum_{t=1}^T \tilde{u}_t \leq \sum_{t=1}^T \log_2(1 + \tilde{u}_t) = \sum_{t=1}^T \log_2 \frac{1 + U_t}{1 + U_{t-1}} = \log_2(1 + U_T).$$

□

E. Proof of **Balanced FTRL**

Space of realization plans Let $\Sigma_{\max}^1 := \{x_{1:h} : x_{1:h} \in \Sigma_{\max}^1\}$ be the set of max-player realization plans. We can notice that this space is a restriction of Σ_{\max}^1 to an affine subspace of Σ_{\max}^1 , defined by the X constraints, for each $h \in [2; H]$,

$$\sum_{a_{h-1} \in A(x_{1:h-1})} x_{1:h}(x_{1:h}; a_{h-1}) = \sum_{a_{h-1} \in A(x_{1:h-1})} x_{1:h-1}(x_{1:h-1}; a_{h-1});$$

with the convention $\mathbb{1}_{:0}(x_0; a_0) = 1$. We thus write $\mathbb{1}_{: \max}^1 = (F + u) \setminus \mathbb{R}_0^{A \times}$, with F a linear subspace of $\mathbb{R}^{A \times}$ and $u \in \mathbb{1}_{: \max}^1$.

Average profile We especially get that the average profile $\mathbb{1}_{: \max}^T$ is also in $\mathbb{1}_{: \max}^1$. We define an average realization plan of the max-player as one of the $\mathbb{1}_{: \max}^1$ that verifies

$$\mathbb{1}_{: \max}^T = \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{: \max}^t$$

Doing the same step for, an average realization of the min-player, we get the existence of an average profile. Note that if the average profile is positive over all information states, then it is unique. It has the same properties as the one defined in [Kozuno et al. \(2021\)](#).

Global regularizer We will assume that the h functions are supported on $\mathbb{R}_0^{A(x_h)}$, and denote by: $\mathbb{R}_0^{A(x_h)} \rightarrow \mathbb{R}$ the function $(\mathbb{1}_{:1}) = \prod_{h=1}^H \frac{1}{p_{1:h}^2} h(\mathbb{1}_{:1:h} p_{1:h}^2)$. This function is strictly convex and non-positive when using either Tsallis or Shannon entropy as h function. Its definition with the coordinates of the realization plan will be important as we will consider its gradient. The associated Bregman divergence is

$$D(\mathbb{1}_{:1}; \mathbb{2}_{:1}) := (\mathbb{1}_{:1}) - (\mathbb{2}_{:1}) + r(\mathbb{2}_{:1}; \mathbb{1}_{:1}, \mathbb{2}_{:1})$$

We first give a lemma that justifies the use of the balanced transitions.

Lemma E.1. For any $\mathbb{2}_{: \min}$, we have

$$\sum_{h=1}^H \sum_{x_h \in X_h} A_{x_h} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)} = A_X$$

Furthermore, when the size of the action set is fixed, we also have for any $\mathbb{2}_{:H}$

$$\sum_{x_h \in X_h} A_{x_h} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)} = A^{h,H} A_X$$

Proof. We first show by induction on h that for all $x_h \in X_h$,

$$\sum_{h^0=1}^{h-1} \sum_{x_{h^0} \in X_{h^0}} A_{x_{h^0}} \frac{p_{1:h^0}(x_{h^0})}{p_{1:h^0}^2(x_{h^0})} + \sum_{x_h \in X_h} A_{x_h} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)} = A_X \tag{4}$$

For $h = 1$, we have

$$\sum_{x_1 \in X_1} A_{x_1} \frac{p_{1:1}(x_1)}{p_{1:1}^2(x_1)} = A_X \sum_{x_1 \in X_1} p_0(x_1) = A_X$$

Assuming the property holds at step h to obtain the property at step $h+1$ we use

$$\begin{aligned} \sum_{x_{h+1} \in X_{h+1}} A_{x_{h+1}} \frac{p_{1:h+1}(x_{h+1})}{p_{1:h+1}^2(x_{h+1})} &= \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \sum_{(\dots; x_h; a_h; x_{h+1})} A_{x_{h+1}} \frac{p_{1:h+1}(x_{h+1})}{p_{1:h}^2(x_h) p_h^2(x_{h+1}; x_h; a_h)} \\ &= \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \sum_{(\dots; x_h; a_h; x_{h+1})} \frac{p_{1:h+1}(x_{h+1})}{p_{1:h}^2(x_h)} \sum_{(\dots; x_h; a_h; x_{h+1}^0)} A_{x_{h+1}^0} \\ &= \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)} \sum_{(\dots; x_h; a_h; x_{h+1}^0)} A_{x_{h+1}^0} \\ &= \sum_{x_h \in X_h} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)} (A_{x_h} + A_{x_h}) \end{aligned}$$

As for any $x_H \in X_H$, $A_{x_H} = A_{x_H}$ (as no state follows a final state), we obtain the first equality using (4) with H .

For the inequality, we can notice that, if the size of the action set is fixed equal to 2, we have for any $h \in [H-1]$:

$$A_{x_h} = \sum_{x^0 \in X} A(x^0) \prod_{a_h \in A(x_h)} \dots \prod_{x_H \in X_H} A_{x_H}$$

$$= \sum_{a_h \in A(x_h)} \sum_{x_H \in X_H} \dots \sum_{x_H \in X_H} A_{x_H} \prod_{a_h \in A(x_h)} A_{x_h}^{H-h-1} = A_{x_h} A_{x_h}^{H-h-1}$$

The inequality $A_{x_h} A_{x_h}^{H-h-1}$ also trivially holds for $h = H$ as explained before. Using again (4), this time with any $h \in [H]$, we obtain

$$A_X = \sum_{h=0}^{H-1} \sum_{x_h \in X_h} A_{x_h} \frac{p_{1:h^0}(x_{h^0})}{p_{1:h^0}^2(x_{h^0})} + \sum_{x_h \in X_h} A_{x_h} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)}$$

$$= \sum_{x_h \in X_h} A_{x_h} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)}$$

$$= \sum_{x_h \in X_h} A_{x_h} \frac{p_{1:h}(x_h)}{p_{1:h}^2(x_h)}$$

Which concludes the proof. □

Lemma E.2. Assume that the size of the action set is fixed equal to 2, then $\mathbb{E} \overline{A_X}$ and we have the inequality:

$$\sum_{h=1}^H A_{(h-H)=2} \leq 2 + \frac{p}{2}$$

Proof. The first inequality is simply obtained with

$$A_X \leq \sum_{h=1}^H A^h \leq \sum_{h=1}^H 2^h = (2^{H+1} - 2) = H^2$$

where the first inequality comes from the fact that, at each time step, the player remembers its 1 past actions, and the last comes from the fact that $(2^{H+1} - 2) = H^2$ is increasing for $H \geq 3$ and is more than 1 for $H = 1; 2; 3$.

For the second inequality we use

$$\sum_{h=1}^H A_{(h-H)=2} \leq \sum_{h=1}^H 2^{(h-H)=2} = \sum_{h=0}^{H-1} A_{h=2} = 1 + (1 - \frac{p}{2}) = 2 + \frac{p}{2}$$

□

We now try to bound the regret. We first decompose it, using the same decomposition as in [Zimmert & Seldin \(2019\)](#) for the expected part.

Lemma E.3. The regret of **Balanced FTRL** can be decomposed into

$$R_T = \underbrace{\sum_{t=1}^T D_{1:t}(\theta^t)}_{\text{BIAS I}} + \max_{y \in \mathcal{Y}} \underbrace{\sum_{t=1}^T D_{1:t}(y; \theta^t)}_{\text{BIAS II}}$$

$$+ \max_{z \in \mathcal{Z}} \underbrace{[\dots]}_{\text{REG}} + \underbrace{\sum_{t=1}^T D_{1:t}(z; r_{1:t}(\theta^t); r_{1:t}(\theta^t))}_{\text{VAR}}$$

Proof. Let $y_{1:}^2$ be some realization plan. For all T , we decompose the instantaneous regret against $y_{1:}^2$ into

$$D_{1:}^t(y_{1:}^2) = h_{1:}^t(y_{1:}^2) + \sum_{i \in I} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} - \frac{D_{1:}^t(y_{1:}^2)}{e^{t-1}} \right) + \sum_{i \in I} \frac{D_{1:}^t(y_{1:}^2)}{e^t} - \frac{D_{1:}^t(y_{1:}^2)}{e^{t-1}}$$

where $h_{1:}^t(y) = \sup_{z_{1:} \in \Sigma_{1:}} h_{1:}^t(z_{1:}; y)$.

Summing the first two terms over t gives the two BIAS terms. For the REG term, summing over t yields, by telescoping,

$$\sum_{t=1}^T h_{1:}^t(y_{1:}^2) - \frac{D_{1:}^T(y_{1:}^2)}{e^T} = (0) - \frac{D_{1:}^T(y_{1:}^2)}{e^T} + \sum_{t=1}^T \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} - \frac{D_{1:}^t(y_{1:}^2)}{e^{t-1}} \right)$$

as $\frac{D_{1:}^t(y_{1:}^2)}{e^t} \geq 0$ for the first inequality, and $\frac{D_{1:}^t(y_{1:}^2)}{e^t} - \frac{D_{1:}^t(y_{1:}^2)}{e^{t-1}}$ is a non-positive function for the second inequality.

We then define $\psi(y) = \sup_{x \in \mathbb{R}^X_0} h(x; y)$, and, thanks to the decomposition

$$\psi(y) = (F + u) \setminus \mathbb{R}^X_0, \text{ get}$$

$$\begin{aligned} D_{1:}^t(y_{1:}^2) &= \sum_{i \in I} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} - \frac{D_{1:}^t(y_{1:}^2)}{e^{t-1}} \right) + \sum_{i \in I} \frac{D_{1:}^t(y_{1:}^2)}{e^t} \\ &= \sum_{i \in I} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} + r_{1:}^t(y_{1:}^2) + g_{1:}^t \right) - \sum_{i \in I} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^{t-1}} + g_{1:}^t \right) \quad (1) \\ &= \sum_{i \in I} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} + r_{1:}^t(y_{1:}^2) \right) - \sum_{i \in I} r_{1:}^t(y_{1:}^2) \quad (2) \\ &= \sum_{i \in I} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} + r_{1:}^t(y_{1:}^2) \right) - \sum_{i \in I} r_{1:}^t(y_{1:}^2) \quad (3) \\ &= \sum_{i \in I} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} + r_{1:}^t(y_{1:}^2) \right) - \sum_{i \in I} r_{1:}^t(y_{1:}^2) \end{aligned}$$

Where we used successively:

(1) As $t = \arg \min_{z_{1:} \in \Sigma_{1:}} \left(\frac{D_{1:}^t(y_{1:}^2)}{e^t} + (z_{1:}) \right)$, we have

$$\frac{D_{1:}^t(y_{1:}^2)}{e^t} + r_{1:}^t(y_{1:}^2) + g_{1:}^t = 0$$

where $g_{1:}^t \in F$.

(2) For all $y \in \mathbb{R}^X$, $(y + g^t) = (y) + h_{1:}^t(y; g^t)$.

(3) Because ψ is a constrained version over \mathbb{R}^X_0 of ψ . And as ψ is convex, $t_{1:} = \arg \max_{x \in \mathbb{R}^X_0} h(x; r_{1:}^t(y_{1:}^2))$ (x), which implies that the maxima associated to $(r_{1:}^t(y_{1:}^2))$ is reached on \mathbb{R}^X_0 . \square

We first give upper bounds on the BIAS terms when using the IX estimations.

Lemma E.4. For any sequence $\epsilon_h = \epsilon > 0$, with probability at least $1 - \epsilon$, we have

$$\text{BIAS I} \leq H^p \frac{1}{2T} + A_X T \quad \text{and} \quad \text{BIAS II} \leq H \frac{1}{2} :$$

If the size of the action set is fixed and $\epsilon_h = A^{-(H-h)/2}$ instead, with probability at least $1 - \epsilon$, we have

$$\text{BIAS I} \leq H^p \frac{1}{2T} + (2 + \frac{1}{2}) A_X T \quad \text{and} \quad \text{BIAS II} \leq (1 + \frac{1}{2}) H :$$

Proof. We can first decompose BIAS I into

$$\text{BIAS I} = \sum_{t=1}^D \sum_{i=1}^t \mathbb{E}[\theta_j | \mathcal{F}_{t-1}] \epsilon + \sum_{t=1}^D \sum_{i=1}^t \mathbb{E}[\theta_j | \mathcal{F}_{t-1}] \epsilon$$

Using Lemma D.2, as $\epsilon = \frac{1}{H}$, we can bound the first term with probability $1 - \epsilon$ by

$$\sum_{t=1}^D \sum_{i=1}^t \mathbb{E}[\theta_j | \mathcal{F}_{t-1}] \epsilon \leq H^p \frac{1}{2T \log(3)} + H^p \frac{1}{2T} :$$

The second term can be upper-bounded by using the Lemma E.1 and for the last equality,

$$\begin{aligned} & \sum_{t=1}^D \sum_{i=1}^t \mathbb{E}[\theta_j | \mathcal{F}_{t-1}] \epsilon \\ &= \sum_{t=1}^D \sum_{h=1}^t \sum_{(x_h; a_h) \in \mathcal{X}_h} p_{1:h}^t(x_h) \frac{1}{p_{1:h}^t(x_h) + p_h^?(x_h; a_h)} \frac{p_{1:h}^t(x_h) p_h^?(x_h; a_h)}{p_{1:h}^t(x_h) + p_h^?(x_h; a_h)} \\ &= \sum_{t=1}^D \sum_{h=1}^t \sum_{(x_h; a_h) \in \mathcal{X}_h} p_{1:h}^t(x_h) p_h^?(x_h; a_h) \\ &= \sum_{t=1}^D \sum_{h=1}^t \sum_{(x_h; a_h) \in \mathcal{X}_h} A_{X_h} \frac{p_{1:h}^t(x_h)}{p_{1:h}^?(x_h)} \\ &= T A_X : \end{aligned}$$

In the setting with a fixed size of the action set, using the depth-wise inequality of Lemma E.1, and Lemma E.2 for the last inequality, we can get instead

$$\begin{aligned} & \sum_{t=1}^D \sum_{i=1}^t \mathbb{E}[\theta_j | \mathcal{F}_{t-1}] \epsilon \\ &= \sum_{t=1}^D \sum_{h=1}^t \sum_{(x_h; a_h) \in \mathcal{X}_h} A_{X_h} \frac{p_{1:h}^t(x_h)}{p_{1:h}^?(x_h)} \\ &= A_X \sum_{t=1}^D \sum_{h=1}^t A^{(H-h)/2} A^{h-H} \\ &= A_X T A^{(H-H)/2} \\ &= (2 + \frac{1}{2}) A_X T : \end{aligned}$$

In order to bound BIAS II, we can use Lemma D.1. Indeed for $(x_h; a_h) \in \mathcal{X}_h$, using $\epsilon_h = \frac{1}{3A_X}$, the constant stopping time $\tau_h(x_h; a_h) = T$ and $\tau_h^0(x_h; a_h) = \tau_h^?(x_h; a_h)$, we have, as $\log(3A_X) \leq H$, with probability at least $1 - \epsilon = 1 - \frac{1}{3A_X}$,

$$\mathbb{E}_h^T(x_h; a_h) - L_h^T(x_h; a_h) \leq \frac{1}{2} \frac{1}{\sqrt{p_{i:h}(x_h)}};$$

By a union bound the inequality holds for all $(x_h; a_h) \in \mathcal{A}(X_h)$ with probability at least $1 - \frac{1}{3}$. In this case, we get the bound, for any $\epsilon \in [0, 1]$, which

$$\begin{aligned} \sum_{t=1}^T D_{1:t}^y(\theta_{1:t}) &= \sum_{h=1}^H \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \frac{1}{2} \frac{1}{\sqrt{p_{i:h}(x_h)}} \\ &= \sum_{h=1}^H \frac{1}{h} \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \frac{1}{\sqrt{p_{i:h}(x_h)}} \\ &= \sum_{h=1}^H \frac{1}{h} \\ &= \frac{1}{2} H \end{aligned}$$

which gives the bound on BIAS II by taking the maximum over $\epsilon \in [0, 1]$. In the second setting we have instead, again with Lemma E.2

$$\sum_{t=1}^T D_{1:t}^y(\theta_{1:t}) \leq \sum_{h=1}^H \frac{1}{h} = \frac{1}{2} \sum_{h=1}^H A^{(h-H)=2} \left(1 + \frac{1}{2}\right)^{-1};$$

□

Depending on the regularizer used in **Balanced FTRL**, we then state some upper bounds of the REG term.

Lemma E.5. Assume that $\epsilon = \frac{1}{2}$, then with Tsallis entropy

$$\text{REG} \leq \frac{H^q}{1-q} A_X^{1-q};$$

and with Shannon entropy

$$\text{REG} \leq H \log(A_X);$$

Furthermore, if the size of the action set is fixed and $\epsilon = \frac{1}{2}$ instead, we have with Shannon entropy

$$\text{REG} \leq \frac{2 + \frac{1}{2}}{2} \log(A_X);$$

Proof. The upper-bound with Tsallis entropy is a consequence of the inequality, as for all $\epsilon \in [0, 1]$,

$$\begin{aligned}
 (1) &= \sum_{h=1}^H \frac{1}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q \\
 &= \sum_{h=1}^H \frac{1}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q \\
 &= \sum_{h=1}^H \frac{1}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q \\
 &= \frac{H}{|A|} \sum_{x \in X} (p_{1:h}^q(x) - p_{1:h}^*(x; a_h))^q
 \end{aligned}$$

When using Shannon entropy, we use that $p_{1:h}^*$ is a probability distribution on $A(X_h)$ for all $h \in [H]$, and get with Jensen inequality

$$\begin{aligned}
 (2) &= \sum_{h=1}^H \frac{1}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q \\
 &= \sum_{h=1}^H \frac{1}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q \\
 &= \sum_{h=1}^H \frac{1}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q \\
 &= \frac{H}{|A|} \sum_{x \in X} (p_{1:h}^q(x) - p_{1:h}^*(x; a_h))^q
 \end{aligned}$$

If the size of the action set is fixed and $|A_h| = |A|^{(H-h)/2}$, we instead get with Lemma E.2

$$(3) \quad \sum_{h=1}^H \frac{1}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q = \sum_{h=1}^H \frac{|A|^{(H-h)/2}}{|A_h|} \sum_{x_h \in X_h} (p_{1:h}^q(x_h) - p_{1:h}^*(x_h; a_h))^q = \frac{2 + \sqrt{2}}{2} \sum_{h=1}^H \log(|A_h|)$$

□

Upper bounding the VAR term is a little harder, but the idea is the same between the two regularizers, with each term of the sum over time being bounded separately in expectation.

Lemma E.6. Let $v_t = D(r_t - r^*)$; $r_t \in \mathcal{R}$; $r^* \in \mathcal{R}$ for all $t \in [T]$. Then with a probability at least $1 - \epsilon$,

$$\text{VAR} \sum_{t=1}^T E[v_t | F_{t-1}] + H \frac{p}{2T}$$

Furthermore, if $\epsilon = \frac{1}{2}$, we have with Tsallis entropy

$$\sum_{t=1}^T E[v_t | F_{t-1}] \leq \frac{T |A| \log(|A|)}{2q(1-q)} H^{1-q}$$

and with Shannon entropy

$$\sum_{t=1}^T E[v_t | F_{t-1}] \leq \frac{T |A|}{2} H$$

If $\mathcal{H} = A^{(\mathcal{H} \setminus h)=2}$ and the size of the action set is fixed A instead, we then have with Shannon entropy

$$\sum_{t=1}^T E[v_t | \mathcal{F}_{t-1}] \leq (1 + \frac{1}{\sqrt{2}}) T A_X :$$

Proof. We can first notice, as \mathcal{H} is supported on $\mathbb{R}^{A \times 0}$ and the estimated losses \mathcal{E}_t are non-negative, that

$$v_t = \sum_{i=1}^D \mathcal{E}_t + \sum_{i=1}^E r(\cdot) \mathcal{E}_t + \sum_{i=1}^D \mathcal{E}_t + \sum_{i=1}^E r(\cdot) \mathcal{E}_t : \mathcal{H} :$$

Which lets us use Lemma D.2 to get, with a probability at least $\frac{1}{3}$,

$$\sum_{t=1}^T v_t \leq E[v_t | \mathcal{F}_{t-1}] \leq \frac{P}{2T \log(3)} \leq \frac{P}{2T} :$$

And obtain the first inequality with the Doob decomposition of the random process $(v_k)_{k=1}^T$, as $\text{VAR} = \sum_{k=1}^T v_k$.

For the upper-bounds of $\sum_{t=1}^T E[v_t | \mathcal{F}_{t-1}]$, we define for all $t \in [T]$, $f_t(u) = \sum_{i=1}^D r(\cdot) u^{\mathcal{E}_t} + \sum_{i=1}^E r(\cdot)$ for $u \in [0, 1]$, such that $f_t(0) = 0$ and $f_t(1) = v_t$. As for both entropy, (\cdot) (respectively $\mathcal{E}_t(y)$) can be decomposed into $\sum_{h=1}^H (x_h; a_h) \mathcal{A}(X_h) \cdot r_{x_h; a_h}(\cdot)$ (respectively $\mathcal{E}_t(y) = \sum_{h=1}^H (x_h; a_h) \mathcal{A}(X_h) \cdot \mathcal{E}_{x_h; a_h}(y(x_h; a_h))$), the derivative of f_t can be expressed with

$$f_t'(u) = \sum_{h=1}^H \sum_{i=1}^A \mathcal{E}_h(x_h; a_h) \cdot \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot r_{x_h; a_h}(\cdot) \cdot u^{\mathcal{E}_h(x_h; a_h) - 1} : (5)$$

When using Tsallis entropy, we have

$$r_{x_h; a_h}(\cdot) = \frac{q}{h} p_{1:h}^q(x_h) \cdot \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\mathcal{E}_h(x_h; a_h) - 1}$$

$$r_{x_h; a_h}(\mathcal{E}_h(y(x_h; a_h))) = \frac{h y(x_h; a_h)}{q p_{1:h}^q(x_h)^q} \cdot u^{\frac{1}{1-q}}$$

which yields

$$\begin{aligned} r_{x_h; a_h}(\mathcal{E}_h(y(x_h; a_h))) &= r_{x_h; a_h}(\cdot) \cdot u^{\mathcal{E}_h(x_h; a_h)} \\ &= \frac{h}{q p_{1:h}^q(x_h)^q} \cdot \frac{q}{h} p_{1:h}^q(x_h) \cdot \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\mathcal{E}_h(x_h; a_h) - 1} \cdot u^{\frac{1}{1-q}} \\ &= \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\frac{h \mathcal{E}_h(x_h; a_h)}{q p_{1:h}^q(x_h)^q} \cdot \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\frac{1}{1-q}} \\ &= \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\frac{h \mathcal{E}_h(x_h; a_h)}{q(1-q) p_{1:h}^q(x_h)^q} \cdot \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\frac{1}{1-q}} \end{aligned}$$

where we use $(1+y)^{\frac{1}{1-q}} = 1 + \frac{y}{1-q}$ for $y \geq 0$. Plugging this in (5), along with $\mathcal{E}_h(x_h; a_h) \cdot \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\mathcal{E}_h(x_h; a_h) - 1} = \mathcal{E}_h(x_h; a_h) \cdot \sum_{i=1}^h r_{x_h; a_h}(\cdot) \cdot u^{\mathcal{E}_h(x_h; a_h) - 1}$, gives

$$f_t^0(u) = \mathbb{E}_{h=1}^{X^t} \mathbb{E}_{(x_h; a_h) \in \mathcal{A}(X_h)} \left[\mathbb{E}_{t_{1:h}(x_h; a_h)} \frac{u^{\theta_h(x_h; a_h)}}{q(1-q)p_{1:h}^2(x_h)^q} t_{1:h}(x_h; a_h)^{1-q} \right]$$

$$= \mathbb{E}_{h=1}^{X^t} \mathbb{E}_{(x_h; a_h) \in \mathcal{A}(X_h)} \left[\mathbb{E}_{t_{1:h}(x_h; a_h)} \frac{u^{\theta_h(x_h; a_h)}}{q(1-q)p_{1:h}^2(x_h)^q} t_{1:h}(x_h; a_h)^{1-q} \right]$$

Summing as u goes from 0 to 1 and conditioning yields

$$\begin{aligned} \mathbb{E}[v_t | F_{t-1}] &= \mathbb{E}_{h=1}^{X^t} \mathbb{E}_{(x_h; a_h) \in \mathcal{A}(X_h)} \left[\mathbb{E}_{t_{1:h}(x_h; a_h)} \frac{1}{2q(1-q)p_{1:h}^2(x_h)^q} t_{1:h}(x_h; a_h)^q | F_{t-1} \right] \\ &= \mathbb{E}_{h=1}^{X^t} \mathbb{E}_{(x_h; a_h) \in \mathcal{A}(X_h)} \left[\frac{1}{2q(1-q)} \frac{p_{1:h}^t(x_h)}{p_{1:h}^2(x_h)^q} t_{1:h}(x_h; a_h)^q \right] \\ &= \frac{1}{2q(1-q)} \mathbb{E}_{h=1}^{X^t} \mathbb{E}_{(x_h; a_h) \in \mathcal{A}(X_h)} \left[\frac{p_{1:h}^t(x_h)}{p_{1:h}^2(x_h)^q} t_{1:h}(x_h; a_h)^q \right] \\ &= \frac{H^{1-q}}{2q(1-q)} A_X^q \end{aligned}$$

where we used $\mathbb{E}[x^q] \leq H^q$, the Hölder inequality, and Lemma E.1 for the last equality.

When using Shannon entropy, we have

$$r_{x_h; a_h}(t_{1:h}(x_h; a_h)) = \frac{p_{1:h}^2(x_h)}{h} [\log(p_{1:h}^2(x_h) t_{1:h}(x_h; a_h)) - 1]$$

$$r_{x_h; a_h}^2(y(x_h; a_h)) = \exp\left(\frac{h}{p_{1:h}^2(x_h)} (y(x_h; a_h) + 1) \log(p_{1:h}^2(x_h))\right)$$

and

$$\begin{aligned} r_{x_h; a_h}^2 &= r_{x_h; a_h}(t_{1:h}(x_h; a_h)) u^{\theta_h(x_h; a_h)} \\ &= \exp\left(\frac{h}{p_{1:h}^2(x_h)} \frac{p_{1:h}^2(x_h)}{h} \log(p_{1:h}^2(x_h) t_{1:h}(x_h; a_h)) u^{\theta_h(x_h; a_h)} \log(p_{1:h}^2(x_h))\right) \\ &= t_{1:h}(x_h; a_h) \exp\left(\frac{h}{p_{1:h}^2(x_h)} \frac{u^{\theta_h(x_h; a_h)}}{h} \log(p_{1:h}^2(x_h))\right) \\ &= t_{1:h}(x_h; a_h) \exp\left(\frac{h}{p_{1:h}^2(x_h)} \frac{u^{\theta_h(x_h; a_h)}}{h} \log(p_{1:h}^2(x_h))\right) : \end{aligned}$$

Using this time $\exp(y) = \sum_{k=0}^{\infty} \frac{y^k}{k!}$ for all $y \geq 0$. With (5), this gives

$$h_t^0(u) = \mathbb{E}_{h=1}^{X^t} \mathbb{E}_{(x_h; a_h) \in \mathcal{A}(X_h)} \left[\mathbb{E}_{t_{1:h}(x_h; a_h)} \frac{u^{\theta_h(x_h; a_h)}}{p_{1:h}^2(x_h)} t_{1:h}(x_h; a_h) \right]$$

$$= \mathbb{E}_{h=1}^{X^t} \mathbb{E}_{(x_h; a_h) \in \mathcal{A}(X_h)} \left[\mathbb{E}_{t_{1:h}(x_h; a_h)} \frac{u^{\theta_h(x_h; a_h)}}{p_{1:h}^2(x_h)} t_{1:h}(x_h; a_h) \right]$$

Again, summing from 0 to 1, using $\eta_h = \frac{1}{2^h}$ and Lemma E.1, leads to

$$\begin{aligned} E[V_t | F_{t-1}] &= \sum_{h=1}^H \sum_{(x_h; a_h) \in 2A(X_h)} \frac{\sum_{i=1}^h p_{1:h}^i(x_h; a_h)}{2p_{1:h}^2(x_h)} \\ &= \frac{1}{2} A_X : \end{aligned}$$

If the size of the action set is fixed and $\eta_h = A^{(H-h)/2}$, we have instead, still with Shannon entropy,

$$\begin{aligned} E[V_t | F_{t-1}] &= \sum_{h=1}^H \sum_{(x_h; a_h) \in 2A(X_h)} \frac{\sum_{i=1}^h p_{1:h}^i(x_h; a_h)}{2p_{1:h}^2(x_h)} \\ &= \frac{1}{2} A_X \sum_{h=1}^H A^{(h-H)/2} \\ &= (1 + \frac{1}{2}) A_X \end{aligned}$$

where we used the depth-wise inequality of Lemma E.1 and Lemma E.2. □

We can now prove the main theorem of Section 4 using the previous lemmas.

Theorem 4.1. Let $\alpha \in (0, 1)$ and $\beta = \log(3A_X T)$. Fix the IX parameters to $\eta_h = \frac{1}{\sqrt{H = (2A_X T)}}$ for all $h \in [H]$.

Using **Balanced FTRL-Tsallis**, $\alpha \in (0, 1)$ and constant learning rates $\eta_h = \frac{1}{\sqrt{2\alpha(1-\alpha)T(H=A_X)^{\alpha-1/2}}}$, with probability at least $1 - \alpha$,

$$R_{\max}^T \leq \frac{1}{2\alpha(1-\alpha)} + 3 \frac{1}{2} \frac{1}{\sqrt{HA_X T}}$$

Using **Balanced FTRL-Shannon** and constant learning rates $\eta_h = \frac{1}{\sqrt{2H \log(A_X) T}} = \frac{1}{\sqrt{2 \log(A_X) T}}$, with probability at least $1 - \alpha$,

$$R_{\max}^T \leq \frac{1}{2 \log(A_X)} + 3 \frac{1}{2} \frac{1}{\sqrt{HA_X T}}$$

Proof. From the decomposition of Lemma E.3, we have

$$R_{\max}^T \leq \text{BIAS I} + \text{BIAS II} + \text{REG} + \text{VAR} :$$

Lemma E.4 directly yields with probability at least $1 - \alpha$

$$\text{BIAS I} + \text{BIAS II} \leq H \frac{1}{2T} + T A_X + H \frac{1}{2}$$

that is minimized with $\eta = \frac{1}{\sqrt{H = (2A_X T)}}$.

Then, using Tsallis entropy we have for the two other terms, from Lemma E.4 and Lemma E.6, with probability at least $1 - \alpha$

$$\text{REG} + \text{VAR} \leq H \frac{1}{2T} + \frac{H^\alpha}{1-\alpha} A_X^{1-\alpha} + \frac{T A_X^\alpha}{2\alpha(1-\alpha)} H^{1-\alpha}$$

minimized with $\eta = \frac{1}{\sqrt{2\alpha(1-\alpha)T H^{\alpha-1/2} A_X^{1-2\alpha}}}$. With the trivial upper bound $H \leq A_X$, we then get by summing the two previous inequalities, with a probability at least

$$R_{\max}^T \leq \frac{1}{3} \frac{1}{\sqrt{2H A_X T}} + \frac{1}{2} \frac{1}{\alpha(1-\alpha)} \frac{1}{\sqrt{HA_X T}}$$

When using Shannon entropy, we instead have (with a probability at least $\frac{1}{3}$) from the same lemmas:

$$\text{REG} + \text{VAR} \leq H \frac{P}{2T} + \frac{H}{2} \log(A_X) + \frac{H}{2} A_X T$$

minimized with $\eta = \frac{P}{2H \log(A_X) (TA_X)}$, which in this case yields by summing, again with probability at least $\frac{1}{3}$,

$$R_{\max}^T \leq 3 \frac{P}{2HTA_X} + \frac{P}{2HTA_X} \log(A_X):$$

□

In the special case of a fixed action set considered in the previous lemma, we can get a better rate for the regret of the max-player, here shown for Shannon entropy.

Theorem 4.4. Assume that the size of the action set is constant equal to A_X and let $\eta \in (0, 1)$, $\beta = \log(3A_X)$. Then, **Balanced FTRL-Shannon** with $\eta = \frac{P}{A_X^{H-h} (2A_X T)}$ and $\beta = \frac{P}{2A_X^{H-h} \log(A_X) (A_X T)}$ yields

$$R_{\max}^T \leq 5 \frac{P}{\log(A_X)} + 8 \frac{P}{A_X T}:$$

As in Corollary 4.3, this allows to compute an optimal profile as long as long as

$$T \geq O((A_X + B_Y) \log(1 + A_X + B_Y)) = \Omega(\dots):$$

Proof. The idea is the same as in the previous theorem, but using instead the upper bounds obtained with a fixed size of the action set that all holds with probability at least $\frac{1}{3}$, as both η and β parameters are in the form $\eta = A_X^{(H-h)/2}$ and $\beta = A_X^{(H-h)/2}$. The case $A_X = 1$ is ignored, as the regret is then trivially 0. With Lemma E.4, we now have:

$$\text{BIAS I} + \text{BIAS II} \leq H \frac{P}{2T} + (2 + \frac{P}{2}) TA_X + (1 + 1 + \frac{P}{2}) \dots$$

minimized with $\eta = \frac{P}{2(2A_X)}$. The REG and VAR terms are in this case upper bounded, using Lemma E.5 and Lemma E.6, by

$$\text{REG} + \text{VAR} \leq H \frac{P}{2T} + \frac{2 + \frac{P}{2}}{2} \log(A_X) + (1 + 1 + \frac{P}{2}) TA_X$$

minimized with $\eta = \frac{P}{2 \log(A_X) (TA_X)}$. By summing everything, we now get, using this time the better inequality $H \leq \frac{P}{A_X}$ from Lemma E.2,

$$R_{\max}^T \leq (2 + 4 \frac{P}{2}) \frac{P}{TA_X} + (2 + 2 \frac{P}{2}) \frac{P}{TA_X} \log(A_X)$$

and we conclude using $\frac{P}{2} \geq 3$ for readability. □

F. Efficient updates and proof of correctness

Dilated entropy equivalence We first state a useful property related to the dilated divergence. For any list of learning rates $\eta := (\eta_h(x_h))_{h,x_h}$, we define the dilated entropy \tilde{r} by

$$\tilde{r}(\cdot) := \sum_{h=1}^H \sum_{(x_h, a_h) \in \mathcal{X}_h \times \mathcal{A}_h} \frac{1}{\eta_h(x_h)} \log \frac{P_{1:h}(x_h; a_h)}{a_h^{0,2A(x_h)} \eta_h(x_h)} :$$

The following lemma is the straight-forward extension of the Lemma 9 by Kozuno et al. (2021) to non-constant learning rates.

Lemma F.1. For any vector η of learning rates and $\eta \in \mathbb{R}^2_{\max}$, the Bregman divergence $D_{\eta}(\cdot; \cdot)$ between the realization plans coincides with the dilated divergence defined in Section 5.

The proof is essentially the same, and is based on the fact (later used in Proposition F.3) that for all positive

$$\eta_{x_h; a_h}(\cdot) = \log(\eta_h(a_h | x_h)) : \quad (6)$$

We then state the result on the equivalence between update (U1) and update (U2).

Proposition F.2. Let $\eta_{t \in [0; T]}$ obtained using **Balanced FTRL** with Shannon entropy and any sequence $\eta_{h \in [H]}$ of learning rates. Then there exists a vector of learning rates and policy $\eta \in \mathbb{R}^2_{\max}$ that can be computed with a time complexity $O(A_X)$ such that it coincides with update (U2),

$$\eta = \arg \min_{\eta \in \mathbb{R}^2_{\max}} D_{\eta}(\cdot; \cdot) + D_{\eta}(\cdot; \cdot) :$$

Proof. From update (U1), $\eta = \arg \min_{\eta \in \mathbb{R}^2_{\max}} D_{\eta}(\cdot; \cdot) + D_{\eta}(\cdot; \cdot)$ where

$$\eta(\cdot) = \sum_{h=1}^H \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \frac{p_{1:h}^2(x_h; a_h)}{2A(X_h)} \log(p_{1:h}^2(x_h; a_h))$$

we first notice that

$$\begin{aligned} \eta(\cdot) &= \sum_{h=1}^H \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \frac{p_{1:h}^2(x_h; a_h)}{h} \log(p_{1:h}^2(x_h; a_h)) \\ &= \sum_{h=1}^H \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \frac{p_{1:h}^2(x_h; a_h)}{h} \log(p_{1:h}(x_h; a_h)) \\ &= \sum_{h=1}^H \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \sum_{h^0=1}^h \frac{p_{1:h}^2(x_h; a_h)}{h} \log(p_{h^0}(a_{h^0} | x_{h^0})) \\ &= \sum_{h^0=1}^H \sum_{(x_{h^0}; a_{h^0}) \in \mathcal{A}(X_{h^0})} \log(p_{h^0}(a_{h^0} | x_{h^0})) \sum_{h=h^0}^H \sum_{(x_h; a_h) \in \mathcal{A}(X_h)} \frac{p_{1:h}^2(x_h; a_h)}{h} \\ &= \sum_{h^0=1}^H \sum_{(x_{h^0}; a_{h^0}) \in \mathcal{A}(X_{h^0})} p_{1:h^0}^2(x_{h^0}; a_{h^0}) \log(p_{h^0}(a_{h^0} | x_{h^0})) \sum_{h=h^0}^H \frac{1}{h} \\ &= \sum_{h^0=1}^H \sum_{(x_{h^0}; a_{h^0}) \in \mathcal{A}(X_{h^0})} \frac{p_{1:h^0}^2(x_{h^0}; a_{h^0})}{h^0} \log(p_{h^0}(a_{h^0} | x_{h^0})) \\ &= \eta^0(\cdot) \end{aligned}$$

where we defined

$$\eta^0(x_{h^0}) := 1 = p_{1:h^0}^2(x_{h^0}) \sum_{h=h^0}^H \frac{1}{h} :$$

This shows that η only differs from η^0 by a linear term of the realization plan: the Bregman divergence of η is thus exactly D_{η^0} . We can then conclude using Lemma F.1 with $\eta = \arg \min_{\eta \in \mathbb{R}^2_{\max}} D_{\eta}(\cdot; \cdot)$

$$\eta(\cdot) = D_{\eta}(\cdot; \cdot) + \eta^0(\cdot) = D_{\eta^0}(\cdot; \cdot) + \eta^0(\cdot) = D_{\eta^0}(\cdot; \cdot) + \eta^0(\cdot) :$$

Algorithm 3 Fast Tree Update

1: Input:
 θ^{t-1} given by update (U2)
 Trajectory $(x_1^{t-1}; a_1^{t-1}; \dots; x_H^{t-1}; a_H^{t-1})$
 Loss θ^{t-1} , old and new learning rates η^{t-1} and η^t along the trajectory
 Base policy π^0

2: Init:
 $Z_{H+1}^t = 1$
 $\pi_h^{t-1} = \pi_h^{t-1}(x_h^{t-1}) = \pi_h^{t-1}(x_h^{t-1})$

3: Iterate:
 For $h = H$ to 1
 $\pi_h^{t-1} = \arg \min_{\pi} \left(\eta^{t-1} \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_{h+1}^t)] + \eta^t \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_h^t)] \right)$
 $Z_h^t = \sum_{a_h \in A(x_h^{t-1})} \pi_h^{t-1}(a_h | x_h^{t-1}) \eta^{t-1} \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_{h+1}^t)] + \sum_{a_h \in A(x_h^{t-1})} \pi_h^{t-1}(a_h | x_h^{t-1}) \eta^t \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_h^t)]$
 For all $a_h \in A(x_h^{t-1})$:
 $\pi_h^t(a_h | x_h^{t-1}) = \frac{\eta^{t-1} \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_{h+1}^t)] + \eta^t \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_h^t)]}{Z_h^t}$
 For all other information sets x_h :
 $\pi_h^t(x_h) = \pi_h^{t-1}(x_h)$

4: Output: π^t

This shows that the two updates minimize two functions that only differ by a constant terms, and are thus equivalent.

The calculation of the π^t is straight-forward with the above formula, while the policy π^t is the minimum over a dilated entropy function that can be recursively computed $\mathbb{E}_{\pi} [r_h(x_h^t; a_h^t)]$, starting from the leaves. \square

We now state Algorithm 3 and a proof of its correctness.

Proposition F.3. Assume that for all $t \in [T]$, the vector η^t of learning rates is only updated along the trajectory of episode t . Let $(\pi^t)_{t \in [T]}$ be the sequence of policies computed with update (U2). Then it coincides with the sequence of policies that would be computed with Algorithm 3.

Proof. We remind that update (U2) define the sequence $(\pi^t)_{t \in [T]}$ with

$$\pi^t = \arg \min_{\pi} \left(\sum_{x_h} \eta^{t-1} \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_{h+1}^t)] + \sum_{x_h} \eta^t \mathbb{E}_{\pi} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_h^t)] \right)$$

Using the decomposition $\frac{1}{\max} = (F + u) \setminus \mathbb{R}^A_{\geq 0}$, Lemma F.1 (on the equivalence between \mathbb{E}_{π} and D_{π}), we then know the existence for all $t \in [T]$ of $g^t \in \mathbb{R}^A_{\geq 0}$ such that

$$r_{\pi^t}(x_h^t) = r_{\pi^{t-1}}(x_h^{t-1}) + \eta^{t-1} \mathbb{E}_{\pi^{t-1}} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_{h+1}^t)] + \eta^t \mathbb{E}_{\pi^t} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_h^t)]$$

Taking the difference on the last equation between t and $t-1$ yields:

$$r_{\pi^t}(x_h^t) - r_{\pi^{t-1}}(x_h^{t-1}) = \eta^{t-1} \mathbb{E}_{\pi^{t-1}} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_{h+1}^t)] + \eta^t \mathbb{E}_{\pi^t} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_h^t)] - \eta^{t-1} \mathbb{E}_{\pi^{t-1}} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_{h+1}^{t-1})] - \eta^{t-1} \mathbb{E}_{\pi^{t-1}} [r_h(x_h^{t-1}; a_h^{t-1}) \log(Z_h^{t-1})] \quad (7)$$

where $\eta^t := g^t - g^{t-1} \in \mathbb{R}^A_{\geq 0}$. And from the normalizing constraints on $\frac{1}{\max}$, π^t is characterized by the vector $\eta^t = (\eta^t(x))_{x \in X}$ with

$$\eta^t(x_h; a_h) = \eta^t(x_h) + \sum_{(x_{h+1}) \in X_{h+1}(x_h; a_h)} \eta^t(x_{h+1})$$

with the sum on the second term being null when $h = H$.

We first show by backward induction on $h \in [H]$ that for all x_h not visited at episode $t-1$, $\eta^t(x_h) = 0$. Indeed, in this case $\eta^{t-1}(x_h) = \eta^t(x_h)$, and the induction property leads to $\eta^t(x_h; a_h) = \eta^t(x_h)$ for all a_h , as the perfect recall assumption

implies any x_{h+1} is also not visited at episode t . For all $a_h \in A(x_h)$, using equation (7) on the component $(x_h; a_h)$ gives with the expression of the gradient of equation (6)

$$\frac{1}{\pi_h^t(x_h)} \log(\pi_h^t(a_h | x_h)) = \frac{1}{\pi_h^t(x_h)} \log(\pi_h^{t-1}(a_h | x_h)) + \nabla_{a_h} \pi_h^t(x_h)$$

which is equivalent to

$$\pi_h^t(a_h | x_h) = \pi_h^{t-1}(a_h | x_h) \exp(\nabla_{a_h} \pi_h^t(x_h) \pi_h^t(x_h))$$

Summing over $a_h \in A(x_h)$ yields $\sum_{a_h} \nabla_{a_h} \pi_h^t(x_h) \pi_h^t(x_h) = 0$ and thus $\nabla_{a_h} \pi_h^t(x_h) = 0$ as $\pi_h^t(x_h) \in (0, 1]$, which concludes the induction. We finally get from the previous equality $\pi_h^t(\cdot | x_h) = \pi_h^{t-1}(\cdot | x_h)$ for all x_h outside of the trajectory at time t , as in Algorithm 3.

For all $h \in [H]$, we are then interested in the visited state x_h^t , and define by convention $\pi_h^t(x_{H+1}^t) = 0$. Then we have for $a_h \in A(x_h^t) \neq a_h^{t-1}$, $\pi_h^t(x_h^t; a_h) = \pi_h^{t-1}(x_h^t; a_h)$ and $\pi_h^t(x_h^t; a_h^t) = \pi_h^{t-1}(x_h^t; a_h^t) + \nabla_{a_h} \pi_h^t(x_h^t)$ using the previous property that π_h^t is null on unvisited states. For all $a_h \in A(x_h^t)$, we now have with equations (6) and (7):

$$\begin{aligned} \frac{1}{\pi_h^t(x_h^t)} \log(\pi_h^t(a_h | x_h^t)) &= \frac{1}{\pi_h^{t-1}(x_h^t)} \log(\pi_h^{t-1}(a_h | x_h^t)) + \frac{1}{\pi_h^t(x_h^t)} - \frac{1}{\pi_h^{t-1}(x_h^t)} \log(\pi_h^0(a_h | x_h^t)) \\ &\quad + \mathbb{1}_{f_{a_h} = a_h^{t-1}} \nabla_{a_h} \pi_h^t(x_h^t) - \nabla_{a_h} \pi_h^{t-1}(x_h^t; a_h^{t-1}) \pi_h^{t-1}(x_h^t) \end{aligned}$$

Multiplying by $\pi_h^t(x_h^t)$ and taking the exponential then leads to:

$$\pi_h^t(a_h | x_h^t) = \pi_h^{t-1}(a_h | x_h^t) \pi_h^{t-1}(\pi_h^0(a_h | x_h^t))^{1 - \pi_h^{t-1}} \exp(\nabla_{a_h} \pi_h^t(x_h^t) \pi_h^t(x_h^t) - \mathbb{1}_{f_{a_h} = a_h^{t-1}} \nabla_{a_h} \pi_h^{t-1}(x_h^t) \pi_h^{t-1}(x_h^t))$$

where $\pi_h^{t-1}(\pi_h^0(a_h | x_h^t))^{1 - \pi_h^{t-1}} = \nabla_{a_h} \pi_h^{t-1}(x_h^t; a_h^{t-1}) \pi_h^{t-1}(x_h^t)$ and $\pi_h^{t-1} = \pi_h^t(x_h^t) = \pi_h^{t-1}(x_h^t)$. We again recognize the update of algorithm 3 with $\pi_h^t(x_h^t) = \exp(\nabla_{a_h} \pi_h^t(x_h^t) \pi_h^t(x_h^t))$, as summing the previous equation over $a_h \in A(x_h^t)$ yields:

$$Z_h^t = \sum_{a_h \in A(x_h^t)} \pi_h^t(a_h | x_h^t) \pi_h^{t-1}(\pi_h^0(a_h | x_h^t))^{1 - \pi_h^{t-1}} \exp(\nabla_{a_h} \pi_h^t(x_h^t) \pi_h^t(x_h^t) - \mathbb{1}_{f_{a_h} = a_h^{t-1}} \nabla_{a_h} \pi_h^{t-1}(x_h^t) \pi_h^{t-1}(x_h^t))$$

We therefore obtained that policy π_h^t obtained with update (12) coincides with the output of Algorithm 3 both in visited and unvisited information set at episode t , which concludes the proof. \square

G. Proof of Adaptive FTRL

In all this section, we will use the constants $\beta = \log(3A_X T)$ and $\gamma = 1 + \log_2(T)$ of theorem 5.1 for readability.

This first lemma centralizes the necessary inequalities for the analysis of the regret which can all hold with a high probability.

Lemma G.1. We assume $\beta \geq 1 = 2$ and $\gamma \geq 1$. Then, when using Adaptive FTRL, the following inequalities all hold with a probability at least $1 - \frac{\beta}{T}$:

$$\mathbb{P}_{1:t}^T(x_h; a_h) \leq 1 + 2P_{1:t}^T(x_h); \tag{8}$$

$$\mathbb{E}^T(x_h; a_h) \leq L^T(x_h; a_h) \left(1 + \frac{2^\beta}{T} P_{1:t}^T(x_h) \right); \tag{9}$$

$$\sum_{t=1}^T D_{\pi_{1:t}^t}(\pi_{1:t}^t) \leq \mathbb{E} \left[H^p \frac{\beta}{2T^\beta} \right]; \tag{10}$$

We denote by \mathcal{E} the event under which all these inequalities hold.

Proof. Let $(x_h; a_h) \in A(X_h)$. We can first notice that for all $t \leq T$, $1 + \mathbb{P}_{1:t}^t(x_h; a_h) \leq 2^t$. Indeed, $1 + \mathbb{P}_a^0 = 1$ and recursively, if $1 + \mathbb{P}_{1:h}^t(x_h; a_h) \leq 2^t$, then $\mathbb{P}_{1:h}^{t+1}(x_h; a_h) = (1 + \mathbb{P}_{1:h}^t(x_h; a_h)) \pi_h^t(x_h; a_h) \leq 2^t \pi_h^t(x_h; a_h) \leq 2^{t+1}$, which implies $1 + \mathbb{P}_{1:h}^{t+1}(x_h; a_h) \leq 2^{t+1}$.

In particular, we get $\mathbb{E}_h^T(x_h; a_h) = 2^T$, and define for all $i \in [T]$, $\epsilon_i^0 = \epsilon^i$.

Let $\tau_h^i(x_h; a_h) = \max_{i \in [T]} \text{ft}_T; \tau_h^i(x_h; a_h) \leq \epsilon_i^0$. As $\tau_h^i(x_h; a_h)$ is a decreasing predictable sequence for $i \in [T]$, $\tau_h^i(x_h; a_h)$ is a stopping time, which lets us apply Lemma D.1 with $\tau_h^i(x_h; a_h) = \tau_h^i(x_h; a_h)$, $\epsilon_i^0 = \epsilon^i(3A_X T)$ and $\epsilon_h^0(x_h; a_h) = \epsilon_i^0$ to get with a probability at least $1 - \epsilon_i^0$,

$$\mathbb{E}_h^i(x_h; a_h) - L_h^i(x_h; a_h) \leq \frac{\epsilon_i^0}{2}$$

where $\epsilon_i^0 = \log(3A_X T)$. We especially get that this inequality holds for $(x_h; a_h) \in A(X_h)$ and $i \in [T]$ with a probability at least $1 - \epsilon_i^0$. As $\epsilon_i^0 = 2^{-T}$, we know the existence of $\tau_h^i(x_h; a_h)$, such that $\tau_h^i(x_h; a_h) \leq \tau_h^0(x_h; a_h)$, and the previous inequality used with $\tau_h^i(x_h; a_h)$ then gives:

$$\mathbb{E}_h^T(x_h; a_h) - L_h^T(x_h; a_h) \leq \frac{\epsilon^0}{2} + \frac{\epsilon^0}{\tau_h^T(x_h; a_h)}$$

As the definition of $\tau_h^i(x_h; a_h)$ implies $\tau_h^i(x_h; a_h) \leq T$

The same exact steps (along with Lemma D.1) can also be performed with $\tau_h^i(x_h; a_h)$, and we similarly obtain that with probability at least $1 - \epsilon_i^0$, for all $(x_h; a_h) \in A(X_h)$,

$$\mathbb{P}_{1:h}^T(x_h; a_h) - P_{1:h}^T(x_h; a_h) \leq \frac{\epsilon^0}{\tau_h^T(x_h; a_h)}$$

which gives, using the definition of $\tau_h^T(x_h; a_h)$,

$$\mathbb{P}_{1:h}^T(x_h; a_h) - P_{1:h}^T(x_h; a_h) \leq \frac{\epsilon^0}{\tau_h^T(x_h; a_h)} + \mathbb{P}_{1:h}^T(x_h; a_h) \leq \frac{\epsilon^0}{\tau_h^T(x_h; a_h)} + \mathbb{P}_{1:h}^T(x_h; a_h)$$

and, as $\epsilon^0 = 2^{-T}$ by assumption,

$$\mathbb{P}_{1:h}^T(x_h; a_h) \leq \frac{2^{-T}}{\tau_h^T(x_h; a_h)} + 2P_{1:h}^T(x_h; a_h) \leq 1 + 2P_{1:h}^T(x_h; a_h)$$

and we recognize inequality (8). Plugging this back in the previous inequality also yields

$$\mathbb{E}_h^T(x_h; a_h) - L_h^T(x_h; a_h) \leq \frac{\epsilon^0}{\tau_h^T(x_h; a_h)} + \mathbb{P}_{1:h}^T(x_h; a_h) \leq \frac{\epsilon^0}{\tau_h^T(x_h; a_h)} + 1 + 2P_{1:h}^T(x_h; a_h)$$

and we now recognize inequality (9).

The last inequality is a consequence of the Lemma D.2 used on $\tau_h^i(x_h; a_h)$. Indeed, as $\tau_h^i(x_h; a_h)$ is non-negative and $\mathbb{E}_h^i(\tau_h^i(x_h; a_h)) \leq \epsilon_i^0$, the lemma gives with a probability at least $1 - \epsilon_i^0$

$$\mathbb{E}_h^i(\tau_h^i(x_h; a_h)) \leq \epsilon_i^0 + \mathbb{E}_h^i(\tau_h^i(x_h; a_h)) \leq \epsilon_i^0 + \mathbb{E}_h^i(\tau_h^i(x_h; a_h))$$

As the three inequalities each hold with a probability at least $1 - \epsilon_i^0$, they all hold at the same time with a probability at least $1 - \epsilon_i^0$, which concludes the proof. □

As in Section E, we decompose the regret between the BIAS, REG and VAR terms. Note that the analysis of REG and VAR terms are different, although conceptually the same. The proof is directly based on the algorithmic updates which the sequence $(Z_h^t)_{h \in [H]}$ are defined.

Lemma G.2. Using Adaptive FTRL yields

$$R_{\max}^T = \underbrace{\sum_{t=1}^T D_{1:t}^t(\theta)}_{\text{BIAS I}} + \max_{y: 2} \underbrace{\sum_{t=1}^T D_{1:t}^t(y; \theta)}_{\text{BIAS II}} + \max_{y: 2} \underbrace{D_{\tau+1}^t(y; 0)}_{\text{REG}} + \underbrace{\sum_{t=1}^T D_{1:t}^t(\theta)}_{\text{VAR}} + \frac{1}{t+1} \log Z_1^t$$

Proof. Let $t \in [T]$ and $h \in [H]$. From Algorithm 3, we have for all $a_h \in A(x_h^t)$,

$$Z_h^{t+1}(x_h^t; a_h) = Z_h^t(a_h | x_h^t) \exp \left(\frac{1}{Z_h^t(a_h | x_h^t)} \log(Z_h^t) \right) \mathbb{1}_{a_h = a_h^t} Z_h^{t+1}(x_h^t)^{-t}$$

where $Z_h^t = \sum_{a_h \in A(x_h^t)} \frac{\log Z_{h+1}^t}{Z_h^t}$. We then get, for all $y \in 2^{\max}$,

$$\sum_{a_h \in A(x_h^t)} \sum_{1:h}^y (x_h^t; a_h) \log(Z_h^{t+1}(a_h | x_h^t)) = \sum_{a_h \in A(x_h^t)} \sum_{1:h}^y (x_h^t; a_h) \log(Z_h^t(a_h | x_h^t)) + (1 - \sum_{1:h}^y (x_h^t; a_h)) \log(Z_h^t(a_h^t | x_h^t)) + \log(Z_h^t)$$

$$\sum_{1:h}^y (x_h^t; a_h) \log(Z_h^{t+1}(a_h^t | x_h^t)) + \frac{1}{\sum_{1:h}^y (x_h^t; a_h)} \log(Z_{h+1}^t) :$$

Using this equality then gives, with the convention $\sum_{1:h}^y(\cdot) = 1$,

$$D_{1:t}^t(y; \theta) = \sum_{h=1}^t \sum_{a_h \in A(x_h^t)} \sum_{1:h}^y (x_h^t; a_h) \log(Z_h^{t+1}(a_h | x_h^t)) - \sum_{h=1}^t \sum_{a_h \in A(x_h^t)} \sum_{1:h}^y (x_h^t; a_h) \log(Z_h^t(a_h | x_h^t)) + (1 - \sum_{1:h}^y (x_h^t; a_h)) \log(Z_h^t(a_h^t | x_h^t)) + \log(Z_h^t)$$

$$= \sum_{h=1}^t \sum_{a_h \in A(x_h^t)} \sum_{1:h}^y (x_h^t; a_h) \log \frac{Z_h^{t+1}(a_h | x_h^t)}{Z_h^t(a_h | x_h^t)} + \sum_{h=1}^t \log \frac{Z_h^{t+1}(a_h^t | x_h^t)}{Z_h^t(a_h^t | x_h^t)} + (1 - \sum_{1:h}^y (x_h^t; a_h)) \log \frac{Z_h^t(a_h^t | x_h^t)}{Z_h^t(a_h^t | x_h^t)} + \log(Z_h^t)$$

$$= \sum_{h=1}^t \sum_{a_h \in A(x_h^t)} \sum_{1:h}^y (x_h^t; a_h) \frac{1}{Z_h^t(a_h | x_h^t)} \log \frac{Z_h^{t+1}(a_h | x_h^t)}{Z_h^t(a_h | x_h^t)} + \frac{1}{Z_h^t(a_h^t | x_h^t)} \log \frac{Z_h^{t+1}(a_h^t | x_h^t)}{Z_h^t(a_h^t | x_h^t)} + \frac{1}{Z_h^t(a_h^t | x_h^t)} \log(Z_h^t)$$

$$+ \frac{1}{Z_h^t(a_h^t | x_h^t)} \log \frac{Z_h^{t+1}(a_h^t | x_h^t)}{Z_h^t(a_h^t | x_h^t)} :$$

Then, using that $Z_h^t(x_h)$ only updates when $a_h = x_h^t$ and in the second equality telescoping, with $\sum_{1:h}^y = 0$, we get

$$\begin{aligned} \sum_{t=1}^T D_{1:t}^y(\theta) &= \sum_{t=1}^T \sum_{h=1}^H \frac{1}{2A(X_h)} \sum_{a_h} p_{1:h}^t(x_h; a_h) \left[\frac{1}{h^{t+1}} \log \frac{y_h(a_h; x_h)}{h^{t+1}(a_h; x_h)} + \frac{1}{h^t(x_h)} \log \frac{y_h(a_h; x_h)}{h^t(a_h; x_h)} \right] \\ &\quad + \frac{1}{h^{t+1}(x_h)} \frac{1}{h^t} \log \frac{y_h^a}{h^t(x_h; a_h)} + \sum_{t=1}^T \frac{1}{h^{t+1}(x_1^t)} \log(Z_1^t) \\ &= D_{1:T+1}^y; T+1 + D_{1:T+1}^y; 0 + \sum_{t=1}^T \frac{1}{h^{t+1}(x_1^t)} \log(Z_1^t) \\ &= D_{1:T+1}^y; 0 + \sum_{t=1}^T \frac{1}{h^{t+1}(x_1^t)} \log(Z_1^t) \end{aligned}$$

The rest of the lemma then follows by the decomposition

$$D_{1:t}^y; \theta = D_{1:t}^y; \theta + D_{1:t}^y; \theta + D_{1:t}^y; \theta + D_{1:t}^y; \theta$$

and the expression of the regret as a scalar product. □

The following lemma gives a bound with high probability on the two BIAS terms for Adaptive FTRL.

Lemma G.3. Under even E , assuming $\epsilon \leq 1$, we have

$$\text{BIAS I} \leq H \sqrt{\frac{P}{2T}} + A_X; \quad \text{BIAS II} \leq X + \frac{2}{\epsilon} HT$$

Proof. We first decompose BIAS I into two terms,

$$\sum_{t=1}^T D_{1:t}^y; \theta = \sum_{t=1}^T D_{1:t}^y; \theta + \sum_{t=1}^T D_{1:t}^y; \theta$$

The first term is directly upper-bounded by $\sqrt{\frac{P}{2T}}$ using the inequality (10). For the second term, we have

$$\begin{aligned} \sum_{t=1}^T D_{1:t}^y; \theta &= \sum_{t=1}^T \sum_{h=1}^H \frac{1}{2A(X_h)} \sum_{a_h} p_{1:h}^t(x_h; a_h) \left[\frac{1}{h^t(x_h; a_h)} - \frac{1}{h^t(x_h; a_h) + h^t(x_h; a_h)} \right] \\ &= \sum_{t=1}^T \sum_{h=1}^H \frac{1}{2A(X_h)} \sum_{a_h} p_{1:h}^t(x_h; a_h) \frac{h^t(x_h; a_h)}{h^t(x_h; a_h) + h^t(x_h; a_h)} \\ &= \sum_{h=1}^H \sum_{a_h} \frac{p_{1:h}^T(x_h; a_h)}{1 + p_{1:h}^T(x_h; a_h)} \\ &\leq \sum_{h=1}^H \sum_{a_h} \log_2(1 + p_{1:h}^T(x_h; a_h)) \\ &\leq \sum_{h=1}^H \sum_{a_h} \log_2(2 + 2 p_{1:h}^T(x_h; a_h)) \\ &\leq A_X \end{aligned}$$

as $\beta = 1 + \log_2(1 + T)$ by definition, where we used Lemma D.3 for β , as $\mathbb{P}_{1:h}^t(x_h; a_h) \leq \frac{1}{\beta}(1 + \mathbb{P}_{1:h}^{t-1}(x_h; a_h))$ and inequality (8) for β .

In order to bound BIAS II, we use inequality (9) and get for all $t \geq \max$,

$$D_{1:t}^y \leq \mathbb{E} \sum_{h=1}^H \sum_{x_h} \sum_{a_h} \frac{1}{2A(x_h)} \left(y_{1:h}^t(x_h; a_h) + \frac{2^0 D_{1:t}^y}{\mathbb{P}_{1:t}^T} \mathbb{E} X + \frac{2^0}{HT} \right)$$

□

Contrary to the analysis of the first algorithm, the bound on the REG term now only holds with high probability because of the adaptive learning rate.

Lemma G.4. Under the event \mathcal{E} , we have the bound

$$\text{REG} \leq \frac{2H \log(A_X)}{X + T}$$

Proof. Using inequality (8) and the fact that ρ is the uniform distribution, we get for any $t \geq \max$

$$\begin{aligned} D_{T+1}^y &\leq \sum_{h=1}^H \sum_{x_h} \sum_{a_h} \frac{1}{2A(x_h)} \frac{1}{T+1} \sum_{x_h} y_{1:h}^t(x_h; a_h) \log \frac{y_{1:h}^t(a_h|x_h)}{\rho_{1:h}^t(a_h|x_h)} \\ &\leq \sum_{h=1}^H \sum_{x_h} \sum_{a_h} \frac{1}{2A(x_h)} \max_{x_{h^0}} \frac{1}{1 + \mathbb{P}_{1:h^0}^T(x_{h^0}^0)} \sum_{x_h} y_{1:h}^t(x_h; a_h) \log(A_X) \\ &\leq \frac{\log(A_X)}{2} \sum_{h=1}^H \sum_{x_h} \sum_{a_h} \max_{x_{h^0}} \left(4 + \frac{1}{A_X} \sum_{a_{h^0}^0} \frac{1}{2A(x_h)} \sum_{x_h} y_{1:h}^t(x_h; a_h) \right) \\ &\leq \frac{\log(A_X)}{2} \sum_{h=1}^H \sum_{x_h} \sum_{a_h} \max_{x_{h^0}} \left(2 + 2\mathbb{P}_{1:h^0}^T(x_{h^0}^0) \sum_{x_h} y_{1:h}^t(x_h; a_h) \right) \\ &\leq \frac{2 \log(A_X)}{2} \sum_{h=1}^H \sum_{x_h} \sum_{a_h} \frac{1}{1 + \mathbb{P}_{1:h}^T(x_h)} \sum_{x_h} y_{1:h}^t(x_h; a_h) \\ &\leq \frac{2H \log(A_X)}{X + T} \end{aligned}$$

and we conclude taking the max over $t \geq \max$.

□

Upper bounding the VAR term is however more difficult. We give a preliminary lemma that requires the assumption that the individual learning rates of the information states at any given time increase along a trajectory. This hypothesis is trivially satisfied by Adaptive FTRL thanks to the particular definition of η_h^t .

Lemma G.5. For all $t \geq [T]$, if the learning rates at time t increase along the trajectory, then

$$D_{1:t}^E \leq \frac{1}{t+1} \log Z_1^t + \frac{H}{2} \sum_{h=1}^H \eta_h^t(x_h^t; a_h^t)$$

Proof. For readability we will define for all $h \geq 2$ in this proof, $\eta_h^t := \eta_h^t(a_h^t|x_h^t)$, $\eta_h^t := \eta_h^t(x_h^t)$, $\eta_h^t := \eta_h^t = \eta_{h+1}^t$ and $\eta_h^t := \eta_h^t(x_h^t; a_h^t)$.

Using Hölder inequality, we first get:

$$\begin{aligned}
 Z_h^t &= \sum_{a_h \in \mathcal{A}(x_h^t)} \mu_h^t(a_h|x_h^t)^{\alpha_h^t} \mu_h^0(a_h|x_h^t)^{1-\alpha_h^t} \exp \left[-\mathbb{1}_{\{a_h=a_h^t\}} \eta_h^{t+1}(x_h^t) \bar{\ell}_h^t \right] \\
 &\leq \left(\sum_{a_h \in \mathcal{A}(x_h^t)} \mu_h^t(a_h|x_h^t) \exp \left[-\mathbb{1}_{\{a_h=a_h^t\}} \eta_h^{t+1}(x_h^t) \bar{\ell}_h^t / \alpha_h^t \right] \right)^{\alpha_h^t} \left(\sum_{a_h \in \mathcal{A}(x_h^t)} \mu_h^0(a_h|x_h^t) \right)^{1-\alpha_h^t} \\
 &= \left(1 - \bar{\mu}_h^t + \bar{\mu}_h^t \exp \left[-\eta_h^{t+1}(x_h^t) \bar{\ell}_h^t / \alpha_h^t \right] \right)^{\alpha_h^t}.
 \end{aligned}$$

This yields as $\bar{\ell}_h^t = \tilde{\ell}_h^t - \log(Z_{h+1}^t) / \eta_{h+1}^{t+1}(x_{h+1}^t)$ and $\alpha_h^t = \eta_h^{t+1}(x_h^t) / \bar{\eta}_h^t$, using convexity and $\kappa_h^t \leq 1$,

$$\begin{aligned}
 (Z_h^t)^{1/\alpha_h^t} &\leq 1 - \bar{\mu}_h^t + \bar{\mu}_h^t (Z_{h+1}^t)^{\bar{\eta}_h^t / \eta_{h+1}^{t+1}(x_{h+1}^t)} \exp \left(-\bar{\eta}_h^t \tilde{\ell}_h^t \right) \\
 &= 1 - \bar{\mu}_h^t + \bar{\mu}_h^t (Z_{h+1}^t)^{\kappa_h^t / \alpha_{h+1}^t} \exp \left(-\bar{\eta}_h^t \tilde{\ell}_h^t \right) \\
 &\leq 1 - \bar{\mu}_h^t + \bar{\mu}_h^t \left(\kappa_h^t (Z_{h+1}^t)^{1/\alpha_{h+1}^t} + (1 - \kappa_h^t) \right) \exp \left[-\bar{\eta}_h^t \tilde{\ell}_h^t \right].
 \end{aligned}$$

We can then define the sequence W_h^t such that $W_{H+1}^t := 1$, and for which the previous relation is always an equality, i.e.

$$W_h^t := 1 - \bar{\mu}_h^t + \bar{\mu}_h^t \left(\kappa_h^t W_{h+1}^t + (1 - \kappa_h^t) \right) \exp \left[-\bar{\eta}_h^t \tilde{\ell}_h^t \right]$$

and we can get by induction, for all $h \in [H]$, $W_h^t \geq (Z_h^t)^{1/\alpha_h^t}$. If we introduce the independent auxiliary variables $y_h^t \sim \mathcal{B}(\kappa_h^t)$ and $z_h^t \sim \mathcal{B}(\bar{\mu}_h^t)$, with the associated products $y_{h:h^0}^t = \prod_{h \leq i \leq h^0} y_i^t$ and $z_{h:h^0}^t = \prod_{h \leq i \leq h^0} z_i^t$, we recursively get

$$W_h^t = \mathbb{E}_{y^t, z^t} \left[\exp \left(- \sum_{h \leq h^0 \leq H} y_{h:h^0}^t z_{h:h^0}^t \bar{\eta}_{h^0}^t \tilde{\ell}_{h^0}^t \right) \right].$$

Finally, $\log(x) \leq -1 + x$, $e^{-x} \leq 1 - x + x^2/2$ and the above relation leads to

$$\begin{aligned}
 \log(Z_1^t) &\leq \alpha_1^t \log(W_1^t) \\
 &= \alpha_1^t \log \left(\mathbb{E}_{y^t, z^t} \left[\exp \left(- \sum_{h=1}^H y_{1:h-1}^t z_{1:h}^t \bar{\eta}_h^t \tilde{\ell}_h^t \right) \right] \right) \\
 &\leq \alpha_1^t \left(-1 + \mathbb{E}_{y^t, z^t} \left[\exp \left(- \sum_{h=1}^H y_{1:h-1}^t z_{1:h}^t \bar{\eta}_h^t \tilde{\ell}_h^t \right) \right] \right) \\
 &\leq -\alpha_1^t \mathbb{E}_{y^t, z^t} \left[\sum_{h=1}^H y_{1:h-1}^t z_{1:h}^t \bar{\eta}_h^t \tilde{\ell}_h^t \right] + \frac{1}{2} \alpha_1^t \mathbb{E}_{y^t, z^t} \left[\left(\sum_{h=1}^H y_{1:h-1}^t z_{1:h}^t \bar{\eta}_h^t \tilde{\ell}_h^t \right)^2 \right] \\
 &\leq -\alpha_1^t \sum_{h=1}^H \bar{\eta}_1^t \bar{\mu}_{1:h}^t \tilde{\ell}_h^t + \frac{1}{2} \alpha_1^t \mathbb{E}_{y^t, z^t} \left[H \sum_{h=1}^H y_{1:h-1}^t z_{1:h}^t \left(\bar{\eta}_h^t \tilde{\ell}_h^t \right)^2 \right] \\
 &= -\eta_1^{t+1}(x_1^t) \sum_{h=1}^H \bar{\mu}_{1:h}^t \tilde{\ell}_h^t + \frac{H}{2} \alpha_1^t \sum_{h=1}^H \bar{\eta}_1^t \bar{\mu}_{1:h}^t \bar{\eta}_h^t \left(\tilde{\ell}_h^t \right)^2 \\
 &\leq -\eta_1^{t+1}(x_1^t) \sum_{h=1}^H \bar{\mu}_{1:h}^t \tilde{\ell}_h^t + \frac{H}{2} \eta_1^{t+1}(x_1^t) \sum_{h=1}^H \bar{\eta}_h^t \tilde{\ell}_h^t
 \end{aligned}$$

where we used $\tilde{\ell}_h^t \mu_{1:h}^t \leq 1$ for the last inequality, and we finally obtain

$$\sum_{h=1}^H \mu_{1:h}^t \tilde{\ell}_h^t + \frac{1}{\eta_1^{t+1}(x_1^t)} \log(Z_1^t) \leq \frac{H}{2} \sum_{h=1}^H \bar{\eta}_h^t \tilde{\ell}_h^t$$

which concludes the proof. \square

We then directly use this lemma on all episodes to upper bound the VAR term.

Lemma G.6. *Under the event \mathcal{E} , assuming $\gamma \geq 1$, we have the bound*

$$\text{VAR} \leq \frac{\eta}{2} H A_{\mathcal{X}} v.$$

Proof. Thanks to Lemma G.5, we have

$$\begin{aligned} \text{VAR} &\leq \frac{H}{2} \sum_{t=1}^T \sum_{h=1}^H \eta_h^t(x_h^t) \tilde{\ell}_h^t(x_h^t, a_h^t) \\ &\leq \frac{H}{2} \sum_{t=1}^T \sum_{h=1}^H \sum_{(x_h, a_h) \in \mathcal{A}(\mathcal{X}_h)} \eta_h^t(x_h) \tilde{p}_{1:h}^t(x_h, a_h) \\ &= \frac{\eta H}{2} \sum_{t=1}^T \sum_{h=1}^H \sum_{x_h \in \mathcal{X}_h} \frac{\tilde{p}_{1:h}^t(x_h, a_h)}{1 + \tilde{P}_{1:h}^{t-1}(x_h)} \\ &= \frac{\eta H}{2} \sum_{h=1}^H \sum_{x_h \in \mathcal{X}_h} \sum_{t=1}^T A_{x_h} \frac{\tilde{p}_{1:h}^t(x_h)}{1 + \tilde{P}_{1:h}^{t-1}(x_h)} \\ &\leq \frac{\eta H}{2} \sum_{h=1}^H \sum_{x_h \in \mathcal{X}_h} A_{x_h} \log_2 \left(1 + \tilde{P}_{1:h}^T(x_h) \right) \\ &= \frac{\eta H}{2} \sum_{h=1}^H \sum_{x_h \in \mathcal{X}_h} A_{x_h} \log_2 \left(1 + \frac{1}{A_{x_h}} \sum_{a_h \in \mathcal{A}(x_h)} \tilde{P}_{1:h}^T(x_h, a_h) \right) \\ &\leq \frac{\eta H}{2} \sum_{h=1}^H \sum_{x_h \in \mathcal{X}_h} A_{x_h} \log_2 \left(1 + \frac{1}{A_{x_h}} \sum_{a_h \in \mathcal{A}(x_h)} (1 + 2P_{1:h}^T(x_h, a_h)) \right) \\ &\leq \frac{\eta H}{2} \sum_{h=1}^H \sum_{x_h \in \mathcal{X}_h} A_{x_h} \log_2 (2 + 2T) \\ &= \frac{\eta v H}{2} A_{\mathcal{X}} \end{aligned}$$

where we used $\tilde{p}_{1:h}^t(x_h^t) \leq 1/(\gamma_h^t(x_h^t, a_h^t) A_{x_h^t}) \leq (1 + \tilde{P}_{1:h}^{T-1}(x_h^t, a_h^t))/(\gamma A_{x_h^t}) \leq 1 + \tilde{P}_{1:h}^{T-1}(x_h^t)$ for the assumption of Lemma D.3, and inequality 8 for the next to last inequality. \square

We now have all the tools needed to upper bound the regret with high probability when using **Adaptive FTRL**.

Theorem 5.1. *Let $\iota' = \log(3A_{\mathcal{X}}T/\delta)$ and $v = 1 + \log_2(1 + T)$. For any $\delta \in (0, 1)$, using **Adaptive FTRL** with $\eta = 2\sqrt{\iota'T/(vA_{\mathcal{X}})}$, $\gamma = \sqrt{2\iota'HT/(vA_{\mathcal{X}})}$ and μ^0 the uniform policy yields with a probability at least $1 - \delta$,*

$$\mathbb{R}_{\max}^T \leq 6H\sqrt{\iota'vA_{\mathcal{X}}T}.$$

As in Corollary 4.3, this leads to the computation of an ε -optimal profile as long as

$$T \geq \mathcal{O} \left(H^2 (A_{\mathcal{X}} + B_{\mathcal{Y}}) \log(1/\delta + T) \log(T)/\varepsilon^2 \right).$$

