see previous proof

But

$$\frac{1}{\eta} \ln\left(\frac{W_{t+1}}{W_t}\right) = \frac{1}{\eta} \ln\left(\sum_{i=1}^{K} P_{i,t}\, e^{-\eta \ell(a_{i,t}, y_t)}\right) \quad (*)$$

③

Notice that: since ℓ is η_0-exp-concave and thus η-exp-concave, then

$$e^{-\eta \ell\left(\sum^{K} P_{i,t} a_{i,t}, y_t\right)} \geq \sum^{K} P_{i,t}\, e^{-\eta \ell(a_{i,t}, y_t)} \quad (**)$$

We take the log and devide by $-\eta$

$$-\ell\Big(\underbrace{\sum^{K} P_{i,t} a_{i,t}}_{\hat{a}_t}, y_t\Big) \geq \frac{1}{\eta} \ln\left(\frac{W_{t+1}}{W_t}\right) \quad \text{from } (*)$$

$$\Rightarrow -\ell(\hat{a}_t, y_t) \geq -\min_{1 \leq i \leq K} L_{i,T} - \frac{\ln K}{\eta} \quad \text{by } (***) \text{ lower bound}$$

$$\Leftrightarrow \ell(\hat{a}_t, y_t) - \min L_{i,T} \leq \frac{\ln K}{\eta}$$

3. The non-convex case

In section 2 we assumed that $\mathcal{D}$ is convex and $\ell(\cdot, y)$ is convex $\forall y \in \mathcal{Y}$. What can we do if these convexity assumption are not satisfied?

Answer: use randomized algorithm!

<u>Algorithm</u>: (Randomized EWA)

- Parameter $\eta > 0$
- At every round $t \geq 1$

  Compute the weight vector $P_t \in \Delta_K$

$$P_{i,t} = \frac{e^{-\eta \sum_{s}^{t-1} \ell(a_{i,s}, y_s)}}{\sum_{j}^{K} e^{-\eta \sum_{s}^{t-1} \ell(a_{j,s}, y_s)}}$$

- Δ $\hat{a}_t = \sum^{K} P_{i,t} a_{i,t}$ $\quad i \in \{1, 2, \dots, K\}$

  is forbidden since we don't know $\hat{a}_t \in \mathcal{D}$ or not since $\mathcal{D}$ is not convex

Instead, pick $I_t \in \{1, \dots, K\}$ at random st $\forall i \in \{1, \dots, K\}$

$\mathbb{P}(I_t = i \mid I_1, \dots, I_{t-1}) = P_{i,t}$

In other words, $\mathcal{L}(I_t \mid I_1, \dots, I_{t-1}) = P_t$

NB: P_t can be random if the statistician has a malicious adversary. In this case $y_t = y_t(I_1, \dots, I_{t-1})$

$a_{\cdot,t} = a_{\cdot,t}(I_1, \dots, I_{t-1})$

$\rightarrow P_t$ is measurable wrt $I_1, \dots, I_{t-2}$

↓ vì P_t phụ c vào $\underbrace{y_{t-1}, a_{t-1}}$ → depend on up to $t-2$

Next we address this more general setting

So now $\hat{a}_t$ will be

$$\hat{a}_t = a_{I_t, t} \quad \text{(we trust expert } I_t \text{ at time } t\text{)}$$

3.1: Bounding the "regret" in expectation

Theorem: Assume that $\ell : \mathcal{D} \times \mathcal{Y} \longrightarrow \mathbb{R}^+$ is bounded in $[B_1, B_2]$, no convexity on $\mathcal{D}$ nor on $\ell(\cdot, y)$ then the randomized EWA algorithm satisfies:

Whatever the adversary (allowed to choose y_t and $a_{\cdot,t}$ as a function of $I_1, \dots, I_{t-1}$)

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell(\hat{a}_t, y_t) - \min_{1 \le i \le K} \sum_{t}^{T} \ell(a_{i,t}, y_t)\right] \le (B_2 - B_1)\sqrt{\frac{T}{2}\ln K}$$

Proof: We use the result of section 2 (no convexity)

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell(\hat{a}_t, y_t)\right] = \sum_{t=1}^{T} \mathbb{E}\left[\mathbb{E}\left(\ell(a_{I_t}, y_t) \mid I_1, \dots, I_{t-1}\right)\right]$$

$y_t, a_{\cdot,t}$ is mesurable wrt $(I_1, \dots, I_{t-1})$

and $\mathbb{P}(I_t = i \mid I_1, \dots, I_{t-1}) = P_{i,t}$

Then the expectation is

$$= \sum^{T} \mathbb{E}\left(\sum_{i=1}^{k} p_{i,t}\, \ell(a_{i,t}, y_t) \right)$$

(Because when we condition on $(I_1, \dots, I_{t-1})$, y_t, $a_{i,t}$ are deterministic, the only randomness comes from I_t)

↳ More formally: if $\begin{cases} Y \text{ is } \sigma(Z)\text{-measurable} \\ \mathcal{L}(X|Z) = \nu(dx|Z) \text{ a.s.} \end{cases}$ then $\mathcal{L}((X,Y)|Z) = \nu(dx|Z) \otimes \delta_Y$ a.s.

As a consequence,

$$\mathbb{E}[\text{regret}(T)] = \mathbb{E}\Bigg[\underbrace{\sum_{t=1}^{T} \sum_{i=1}^{K} p_{i,t}\, \ell(a_{i,t}, y_t) - \min_{1 \le i \le K} \sum_{t=1}^{T} \ell(a_{i,t}, y_t)} \Bigg]$$

We proved at the end of the proof of the 1st thm that this quantity is $\le \frac{\ln K}{\eta} + \frac{\eta T (B_2 - B_1)^2}{8} = (B_2 - B_1)\sqrt{\frac{T}{2}\ln K}$ for $\eta = \frac{1}{B_2 - B_1}\sqrt{\frac{T}{2}\ln K}$

This upper bound is true <u>almost surely</u> (actually: $\forall \omega \in \Omega$).

Therefore, $\mathbb{E}[\text{regret}(T)] \le (B_2 - B_1)\sqrt{\frac{T}{2}\ln K}$ ∎

3.2 : Bounding the regret with high probability

Lemma (Hoeffding - Azuma's inequality)

Let $(X_t)_{t \in \mathbb{N}^*}$ be a $(\mathcal{F}_t)$-adapted processes in L^1 st $X_t \in [A_t, A_t + c_t]$ a.s for some $\mathcal{F}_{t-1}$-measurable r.v A_t and some real number, then for $\forall x > 0$

$$\mathbb{P}\left(\sum_{t=1}^{T} (X_t - \mathbb{E}(X_t | \mathcal{F}_{t-1})) > x \right) \le e^{-\frac{2x^2}{\sum c_t^2}}$$

even better $\forall x \geq 0$,

$$\mathbb{P}\left(\max_{1 \le t \le T} \sum_{s=1}^{t} (X_s - \mathbb{E}(X_s | \mathcal{F}_{s-1})) > x \right) \le e^{-\frac{2x^2}{\sum_{t=1}^{T} c_t^2}}$$

<u>Sketch of the proof</u>

$$Z_t = X_t - \mathbb{E}(X_t | \mathcal{F}_{t-1})$$

Note that $\mathbb{E}(Z_t | \mathcal{F}_{t-1}) = 0$ a.s

$$\mathbb{P}\left(\max_{1 \le t \le T} \sum_{s=1}^{t} Z_s > x \right) = \mathbb{P}\left(\max_{1 \le t \le T} e^{\lambda \sum_{s=1}^{t} Z_s} > e^{\lambda x} \right) \quad \text{for any } \lambda > 0$$

$$\overset{①}{\le} \frac{\mathbb{E}(e^{\lambda \sum^{T} Z_s})}{e^{\lambda x}} \overset{②}{\le} e^{-\lambda x} e^{\frac{\lambda^2}{2} \frac{\sum^{T} c_t^2}{4}} \underset{\text{after optimizing } \lambda}{=} e^{-\frac{2x^2}{\sum c_t^2}}$$

Hint to prove ① and ②: ① is follow from Doob's ineq. for positive submartingale $e^{\lambda \sum^{t} z_s}$

② follows from a proof similar to that of Hoeffding's inq.

Next we use this inequality to bound the regret with high probability.

We prove that a.s

$$\sum_{t}^{T} p_t \cdot \ell_t \leq \min_{1\leq i\leq K} \sum_{t=1}^{T} \ell(a_{i,t}, y_t) + (B_2 - B_1)\sqrt{\frac{T}{2}\ln K}$$

Instead of $\mathbb{E}[\cdot]$ we take $\mathbb{E}[\cdot|\mathcal{F}_{t-1}]$

$$\mathcal{F}_{t-1} = (I_1, \dots, I_{t-1})$$

By using Hoeffding-Azuma's ineq. with $\begin{cases} A_t = B_1 \\ C_t = B_2 - B_1 \end{cases}$

we get that $\forall \delta \in (0,1)$

$$\mathbb{P}\left[\sum^{T}(X_t - \mathbb{E}(X_t|\mathcal{F}_{t-1})) > \sqrt{\frac{\sum^{T} c_t^2}{2}\ln\left(\frac{1}{\delta}\right)}\right] \leq \delta$$

In otherwords, with $\mathbb{P} \geq 1-\delta$ we have

$$\sum^{T}(X_t - \mathbb{E}(X_t|\mathcal{F}_{t-1})) < \sqrt{\frac{\sum c_t^2}{2}\ln\left(\frac{1}{\delta}\right)}$$

$$= (B_2 - B_1)\sqrt{\frac{T}{2}\ln\frac{1}{\delta}}$$

But: $X_t = \ell(\hat{a}_t, y_t)$

$$\mathbb{E}(X_t|\mathcal{F}_{t-1}) = \mathbb{E}(\ell(\hat{a}_t, y_t)|I_1, \dots, I_{t-1})$$

$$= \sum_{i=1}^{K} p_{i,t}\,\ell(a_{i,t}, y_t) = p_t \cdot \ell_t$$

Putting everything together with probability at least $1-\delta$

$$\sum_{t=1}^{T} \ell(\hat{a}_t, y_t) \leq \sum_{t=1}^{T} p_t \cdot \ell_t + (B_2 - B_1)\sqrt{\frac{T}{2}\ln\left(\frac{1}{\delta}\right)}$$

$$\leq \min_{1\leq i\leq K} \sum_{t=1}^{T} \ell(a_{i,t}, y_t) + (B_2 - B_1)\sqrt{\frac{T}{2}\ln K} + (B_2 - B_1)\sqrt{\frac{T}{2}\ln\frac{1}{\delta}}$$

We've just proved the following theorem

Thm Assume ℓ is bounded in $[B_1, B_2]$. Then the randomized EWA algorithm has a regret bounded as follows; whatever the adversary, $\forall \delta \in (0,1)$ with proba $\geqslant 1-\delta$

$$\text{regret}(T) \leq (B_2 - B_1)\sqrt{\frac{T}{2}\ln K} + (B_2 - B_1)\sqrt{\frac{T}{2}\ln\left(\frac{1}{\delta}\right)}$$

Exercise 1/ prove that the randomized EWA algorithm satisfies a.s. $\exists c>0$ st, for T is large enough $T \geqslant T_0(\omega)$

$$\text{regret}(T) \leqslant c\sqrt{T \log T}.$$

2/ $\leqslant c\sqrt{T \ln \ln T}$. (even better)