Exercise 3: Budgeted prediction

Ante-scriptum: we assume in this problem that the horizon T, the budget $m \in \{1, \ldots, T-1\}$ and the loss range $[0, 1]$ are known.

We study a case of prediction of individual sequences when the statistician does not get to see the N–vector of losses at the end of each round, unless she asks for it, which she can only do m times during the T rounds. More formally, the prediction protocol is the following: for all rounds $t = 1, 2, \ldots, T$,

– the statistician picks a distribution $\boldsymbol{p}_t$ over $\{1, \ldots, N\}$ and draws a component I_t at random according to $\boldsymbol{p}_t$;
– simultaneously, the opponent picks a loss vector $(\ell_{1,t}, \ldots, \ell_{N,t}) \in [0, 1]^N$;
– the statistician suffers the loss $\ell_{I_t,t}$ but does not observe it;
– the statistician decides whether she wants to observe the loss vector (and in this case, she observes all of its components); she may only do so if she performed less than $m - 1$ observations so far;
– the opponent observes I_t and $\boldsymbol{p}_t$.

We will construct step by step a strategy for the statistician. We fix a confidence level $\delta \in (0, 1)$.

Random observations and estimated losses (requires Lecture #2)

The statistician will make random decisions about observations. More precisely, she will set $\varepsilon \in (0, 1)$, consider a sequence $Z_1, Z_2, \ldots, Z_T$ of i.i.d. random variables, distributed according to a Bernoulli distribution with parameter ε, and observe the t–th loss vector if and only if $Z_t = 1$.

To abide by the budget constraint, she wants to pick ε such that

$$\mathbb{P}\bigl\{Z_1 + Z_2 + \ldots + Z_T \leqslant m\bigr\} \geqslant 1 - \delta\,.$$

1. Show that $\varepsilon = m/T - (1/T)\sqrt{m/\delta}$ is a suitable choice when $\delta \geqslant 1/m$. You may use Chebychev's inequality to that end.

We define

$$\widehat{\ell}_{j,t} = \frac{\ell_{j,t}}{\varepsilon} Z_t\,.$$

2. Show that for a well-chosen filtration $\mathcal{F} = (\mathcal{F}_t)_{t \geqslant 0}$ to determine, we have

$$\mathbb{E}\Bigl[\widehat{\ell}_{j,t} \,\big|\, \mathcal{F}_{t-1}\Bigr] = \ell_{j,t}\,.$$

Strategy based on these estimated losses (requires Lecture #2)

3. Construct a strategy that never asks for more than m observations and ensures that with probability at least $1 - \delta$,

$$\sum_{t=1}^{T} \sum_{i=1}^{N} p_{i,t} \widehat{\ell}_{i,t} - \min_{j=1,\ldots,N} \sum_{t=1}^{T} \widehat{\ell}_{j,t} \leqslant 2\sqrt{\frac{1}{\varepsilon} \min_{j=1,\ldots,N} \sum_{t=1}^{T} \widehat{\ell}_{j,t} \ln N} + \frac{13}{\varepsilon} \ln N$$

4. Deduce from this a strategy that never asks for more than m observations and whose pseudo-regret

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell_{I_t,t}\right] - \min_{j=1,\ldots,N} \mathbb{E}\left[\sum_{t=1}^{T} \ell_{j,t}\right]$$

is bounded by something of the order of $T\sqrt{(\ln N)/m}$. Please state a precise bound.

Hint: Of course you will take expectations in the bound of Question 3. But there are issues to take care of, like tuning δ and ε.

Note: one can show that $T\sqrt{(\ln N)/m}$ is the optimal order of magnitude of the pseudo-regret; when $m = T$, we are back to the classical case (same setting, same bound) discussed in our series of lectures.