Everything we've talked about has been for IID data. What about adversarial data?

$$\begin{tikzcd} \text{Adv} \arrow[r] \arrow[rrr, bend left, <->] & x_t \arrow[d] & & y_t \arrow[dll, bend left] \\ & \text{Alg} \arrow[r] & p_t \arrow[ull, bend right] & \end{tikzcd}$$

$\pi = ((x_1, p_1, y_1) \ldots (x_t, y_t, p_t))$

π is α MMC if

$$\frac{1}{T}\left|\sum p_t - y_t\right| \le \alpha$$

p α MQC for q if

$$\left[\frac{1}{T}\sum \mathbb{1}[p_t \le y_t] - q\right] \le \alpha$$

ese seem hard to show for adversarial π

...it

Alg 1

For $t=0$ to T
Let x_t be observed
Predict $y_t = 1$

$\Leftarrow$ This is $\frac{1}{T}$-MMC

Proof

Si

$\frac{1}{T}|\Sigma$

$\frac{1}{T}$

I

For a more useful algorithm (though we won't show it yet) it'll help to consider the following.

Suppose A is α-MQC for a quantile q against any adversary. Then it can be used to make an iid α'-MQC by taking a model f $\hat{p}$

$(\hat{x}_t, \hat{y}_t) \sim D$

Give A $(x_t, y_t) = (\hat{x}_t, \hat{y}_t - f(\hat{x}))$

p predict $p_t = A(\pi_t)$

Then $\Delta \sim_{UAR} [p_1, \ldots p_T] \rightarrow \hat{f}(x) = f(x) + \Delta$

is $\alpha' = \alpha + \sqrt{2\log\frac{3}{\delta}}$ MQC

Proof) Since A is α MQC on its transcript

$$\frac{1}{T}\left|\sum \mathbb{1}[(y_t) \le p_t] - q\right| \le \alpha$$

$$\frac{1}{T}\left|\sum \mathbb{1}(\hat{y}_t - f(\hat{x}_t) \le p_t) - q\right| \le \alpha$$

Is it true that $\hat{y}_t - f(\hat{x}_t) \le p_t$ $\approx q$ frac of the time
on avg

Well, p_t's are no longer iid $\quad = A(\Pi_{t-1}, \hat{x}_t)$
and neither are the indicator RVs

So we need a new concentration inequality

Let $X_1, \ldots, X_n$ be RVs (not nec IID!)
$|X_i| \le c_i$. Define $X_{<i} = (X_1, \ldots X_{i-1})$
Then $\forall t > 0$

$$\Pr\left[\left|\sum X_i - \mathbb{E}[X_i \mid X_{<i}]\right| \ge t\right] \le 2\exp\left(\frac{-t^2}{2\sum c_i^2}\right)$$

Azuma's

$\quad$ f of $x_1 \ldots x_t$, $y_1 \ldots y_{t-1}$

This applies to $\frac{1}{T}\mathbb{1}(y_t \le p_t) \perp y_t$

RVs btwn $-\frac{1}{T}, \frac{1}{T}$

Exercise
$E[\mathbb{1}(y_t \le p_t) \mid y_{<t}] = F(p_t)$

So Azuma's $\Rightarrow$

$$\Pr\left[\left|\underbrace{\frac{1}{T}\sum \mathbb{1}(y_t \le p_t)}_{\approx q + \alpha} - \frac{1}{T}\sum F(p_t)\right| \ge \sqrt{\frac{2\ln\frac{2}{\delta}}{T}}\right] \le \delta$$

So $|\Pr[y \le f(x) + \Delta] - q| = |\frac{1}{T} \sum \Pr[y \le f(x) + p_t] - q|$

$|\frac{1}{T} \sum F(p_t) - q|$

$\le \alpha + \sqrt{\frac{2 \ln \frac{2}{\delta}}{T}}$

w.p. $1-\delta$

How do we make an adversarial α-MQC alg?

~~Alg~~ Online MQC (q, η) ← learning rate

Let $p_1 = 0$

for $t = 1 \ldots T$

Observe & ignore x_t

Predict p_t

Observe y_t

$p_{t+1} = p_t + \eta(q - \mathbb{1}(y_t \le p_t))$

Thm: for any sequence of len T, q, & η, ↑ is α MQC for $\alpha \le \frac{1+\eta}{\eta T}$.

$\boxed{\text{Pf}}$

$p_{t+1} = p_t + \eta(q - \mathbb{1}(y_t \le p_t))$

$\Rightarrow \mathbb{1}(y_t \le p_t) = q - \left(\frac{p_{t+1} - p_t}{\eta}\right)$

So $\frac{1}{T}\sum \mathbb{1}[y_t \le p_t] = q - \frac{1}{\eta T}\sum (p_{t+1} - p_t)$

$$= q - \frac{1}{\eta T}[p_{T+1} - p_1]$$

$$= q - \frac{1}{\eta T}[p_{T+1}]$$

Do we have a bound on p_{T+1}?

$\cancel{p_0 = 0},$

$p_{t+1} = p_t + \eta(q - \mathbb{1}(\,))$

$|p_t - p_{t+1}| \le \eta$

If $p_t \ge 1$, $\mathbb{1}[\,] = 1$, $p_{t+1} < p_t$

~~If $p_t \le 0$~~ ~~$p_{t+1} > p_t$~~

$p_{t+1} \le p_t$ $\Rightarrow p_t - p_{t+1} \le \eta$

[Argue to yourselves $\eta \le p_t \le 1+\eta$]

An even dumber algo works

w.p. q Predict $p_t = 1$

w.p. $1-q$ Predict $p_t = 0$

This is $\frac{1}{T}$ MQC...