Overcoming Brittleness in Pareto-Optimal Learning Augmented Algor.

Spyros Angelopoulos
Christoph Dürr
Alex Elenter
Yanni Lefki

One way Trading

1€ → \$

day i exchange rate $p_i \in [1,M]$
exchange fraction $x_i \geq 0$
last day (worst case $p_n = 1$)
exchange remaining $\sum x_i = 1$

profit $= \sum p_i x_i$

OPT $= p^* := \max p_i$
prediction $\hat{p}$

Competitive ratio $= \sqrt{M}$

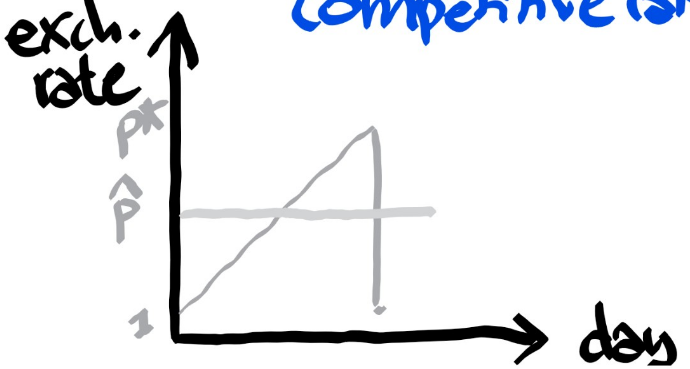

Blindly trust prediction

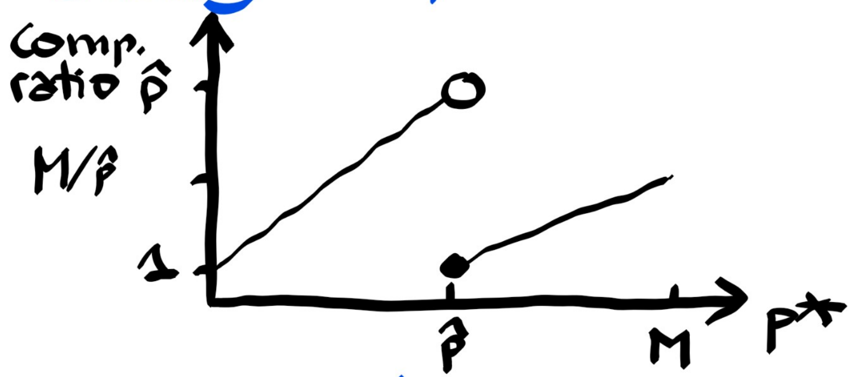

Our approach

Given some profile, is there a strategy which competitive ratio matches the profile?

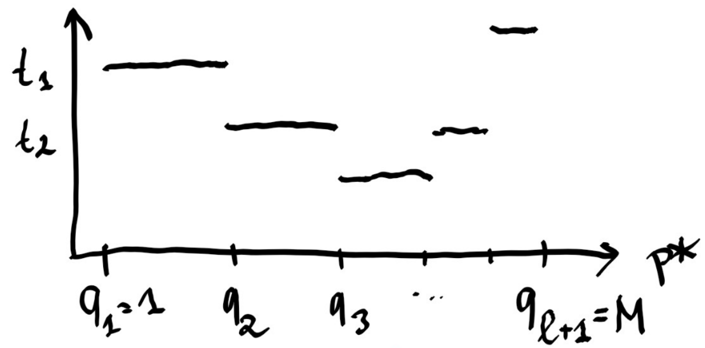

We can answer this question in linear time $O(\ell)$

Continuous time variant ($p_i = i$)

X s.t. $\int_1^M x_i \, di = 1$ $\xrightarrow{\text{easy}}$ competitive ratio profile
$\xleftarrow[\text{hard}]{}$
solve second order differential equation

THE TECHNICAL PART

continous time setting
time = exchange rate $\in [1, M]$

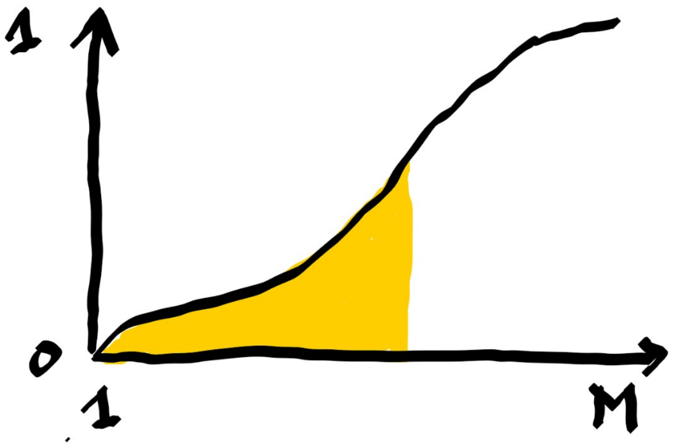

policy = increasing function $[1,M] \to [0,1]$
$f(t)$ = how much we already exchanged by time t

profit so far = $\int_1^t f(x)\,dx$

profit if game stops = $\int_1^t f(x)\,dx + \underbrace{(1-f(x))}_{\text{exchange remaining at rate 1}}$

To guarantee competitive ratio $R > 1$:

$\forall t \in [1,M]$: $\int_1^t f(x)\,dx + (1-f(x)) \underbrace{\geq}_{\text{best policy has} =} \frac{t}{R}$

Optimal f can be found. (Now we understand function given in literature)

Out of reach if R depends on t.
But piecewise constant $R(t)$ is doable

Adaptive Policy

actual problem = discrete time

We end up with more profit, and can aford to wait for better rates