Constant Regret, Generalized Mixability, and Mirror Descent

Zakaria Mhammedi

&

Robert C. Williamson

Australian National University

Motivation

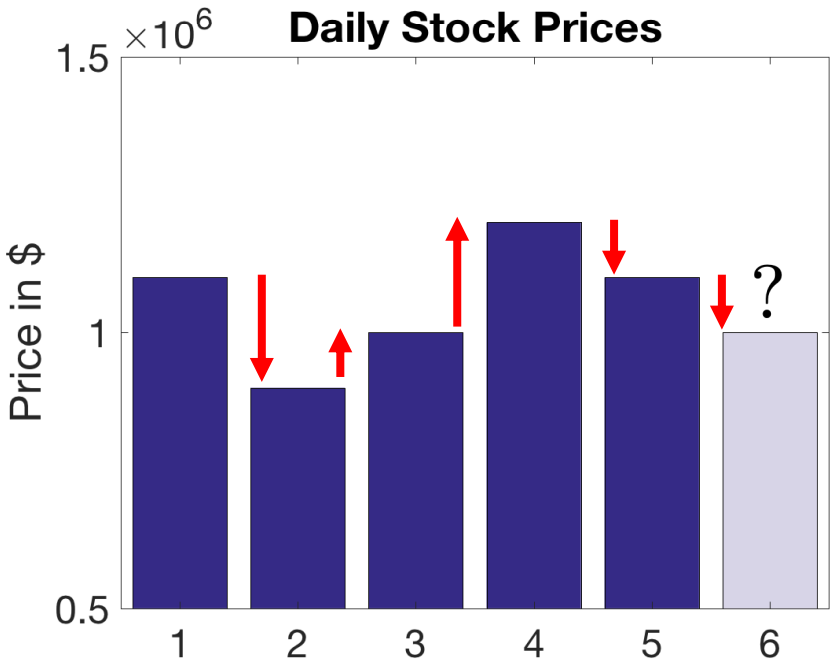

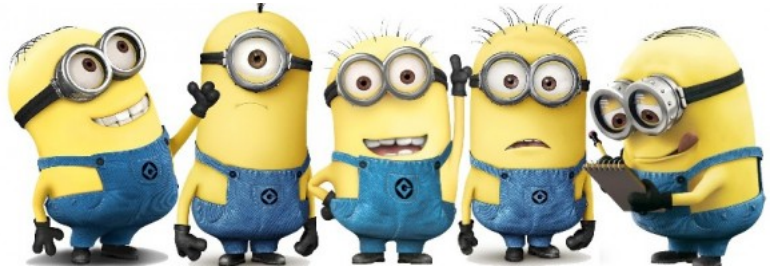

$\mathbf{p}_1$ $\mathbf{p}_2$ $\mathbf{p}_3$ $\mathbf{p}_4$ $\mathbf{p}_5$

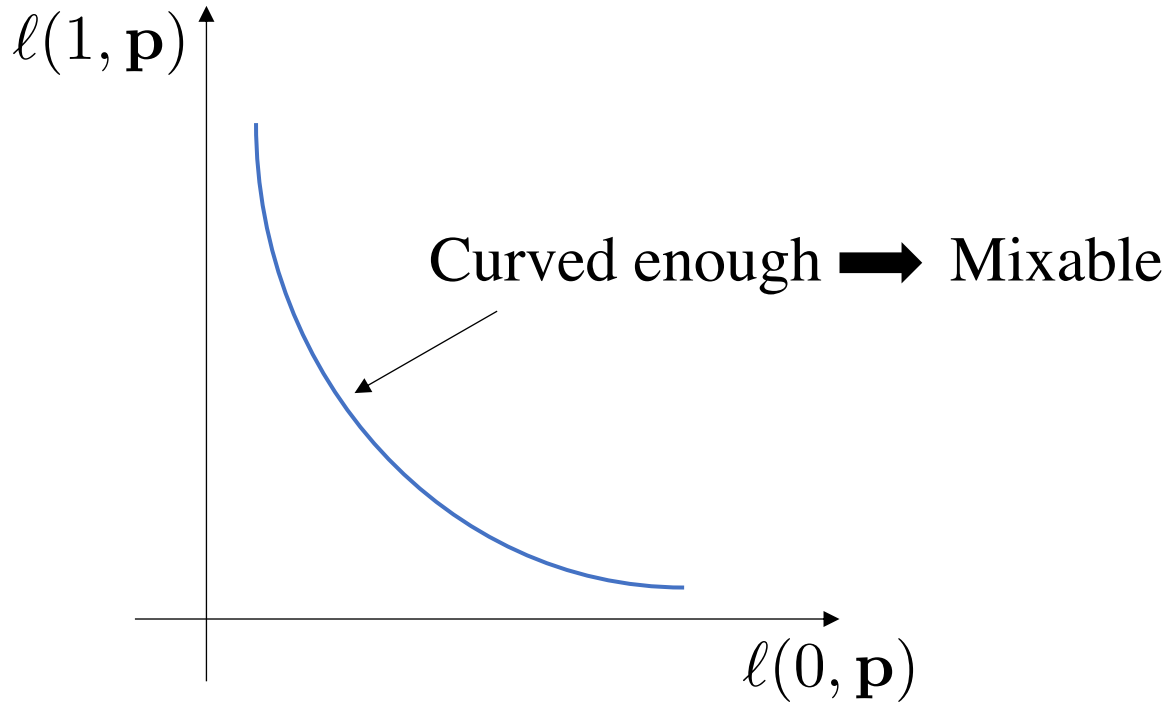

$\text{Mix}(\mathbf{p}_1, \ldots, \mathbf{p}_5)$

Prediction with Expert Advice (PwEA)

At each round $t \in \{1, \ldots, T\}$ the learner;

1. Observes experts' predictions $\boldsymbol{p}_1^t, \ldots, \boldsymbol{p}_k^t \in \mathcal{P}$;
2. Chooses prediction $\boldsymbol{p}_*^t \in \mathcal{P}$;
3. Observes outcome $x^t \in \Omega$;

- Learner suffers $\ell(x^t, \boldsymbol{p}_*^t)$
- Expert k suffers $\ell(x^t, \boldsymbol{p}_k^t)$

$$\sum_{k=1}^{T} \ell(x^t, \boldsymbol{p}_*^t) - \sum_{k=1}^{T} \ell(x^t, \boldsymbol{p}_k^t) \leq R$$

R independent of T ➡ **a constant regret**!

Mixability and the Aggregating Algorithm

- Mixable loss + the Aggregating Algorithm (AA) ➡ a constant regret (Vovk et al. 1998)

Definition: Mixability is a geometric property of the loss function:

$$\ell \text{ is } \eta\text{-mixable} \iff e^{-\eta \mathscr{S}_\ell} \text{ is convex} \quad \text{(Vovk et al. 2009)}$$

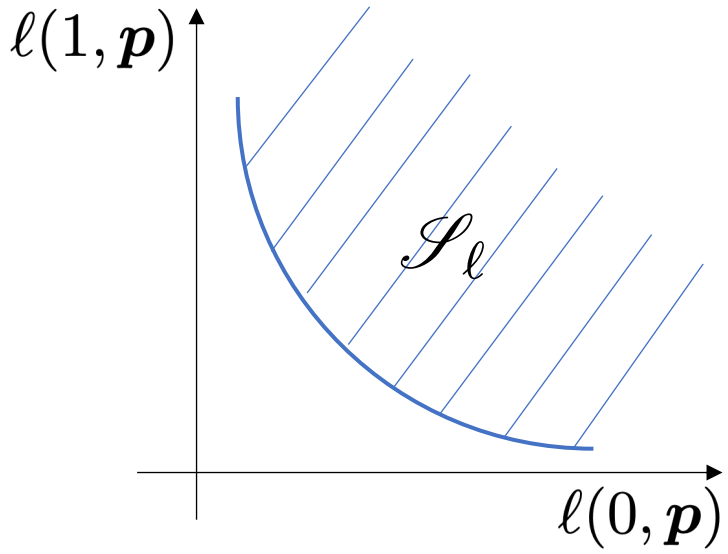

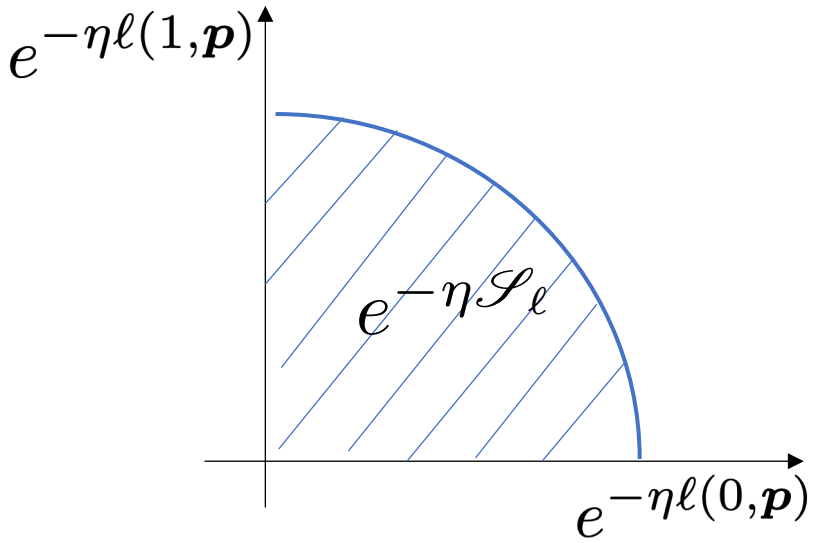

Generalized Mixability + the Generalized AA

- A loss is (η, Φ)-mixable + the generalized AA ➡ a constant regret (Reid et al. 2015)

Definition: The generalized mixability (or (η, Φ)-mixability) condition:

  - For all prior weights $\boldsymbol{q} \in \Delta_K$ on experts;
  - For all experts' predictions $\boldsymbol{p}_1, \ldots, \boldsymbol{p}_K \in \mathcal{P}$;
  - There exists a prediction $\boldsymbol{p}_* \in \mathcal{P}$;

$$\forall x,\ \ell(x, \boldsymbol{p}_*) \leq \inf_{\boldsymbol{\mu} \in \Delta_K} \sum_{k=1}^{K} \mu_k \cdot \ell(x, \boldsymbol{p}_k) + \frac{1}{\eta} D_\Phi(\boldsymbol{\mu}, \boldsymbol{q})$$

Main Results

- ℓ is (η, Φ)-mixable $\iff$ $\frac{\eta_\ell}{\eta}\Phi - \Phi_{\mathrm{sh}}$ is convex on Δ_K
- The Shannon entropy $\Phi_{\mathrm{sh}}(\boldsymbol{q}) \triangleq \boldsymbol{q}^\mathsf{T} \ln \boldsymbol{q}$ is fundamental:
  - ℓ is (η, Φ)-mixable $\implies$ ℓ is $(\eta, \Phi_{\mathrm{sh}})$-mixable
  - The generalized AA (GAA) achieves the lowest worst-case regret using Φ_{sh}

Main Results

- ℓ is (η, Φ)-mixable $\iff$ $\frac{\eta_\ell}{\eta}\Phi - \Phi_{\mathrm{sh}}$ is convex on Δ_K
- The Shannon entropy $\Phi_{\mathrm{sh}}(\boldsymbol{q}) \triangleq \boldsymbol{q}^\top \ln \boldsymbol{q}$ is fundamental:
  - ℓ is (η, Φ)-mixable $\implies$ ℓ is $(\eta, \Phi_{\mathrm{sh}})$-mixable
  - The generalized AA (GAA) achieves the lowest worst-case regret using Φ_{sh}

- We leveraged the similarity between mirror descent and the GAA to devise a new adaptive algorithm (AGAA) for the game of PwEA.
  - The AGAA achieves a negative regret in some cases.

Thank you!

Poster: 10:45 AM -- 12:45 PM @ Room 210 & 230 AB #96