Online Optimization and Games

Lecturer: Jiaming Liang *October 30, 2025*

1 Online linear optimization

The online linear optimization problem can be described as the following game. We are given a closed convex set $S \subseteq \mathbb{R}^n$ (strategy space) and we select some strategy $x_t \in S$ in each round $t \geq 0$. We are then revealed (by the environment or the adversary) a penalty $p_t \in \mathbb{R}^n$. Our goal is to minimize the regret of T rounds of this game:

$$\operatorname{regret}(T) = \sum_{t=0}^{T-1} p_t^\top x_t - \min_{x \in S} \sum_{t=0}^{T-1} p_t^\top x. \tag{1}$$

The online convex optimization problem generalizes online linear optimization in that after choosing a decision $x_t \in S$, the penalty is $f_t(x_t)$ instead of $p_t^\top x_t$, where f_t is convex. The regret of T rounds becomes

$$\operatorname{regret}(T) = \sum_{t=0}^{T-1} f_t(x_t) - \min_{x \in S} \sum_{t=0}^{T-1} f_t(x). \tag{2}$$

We focus the online linear optimization problem in the rest of the lecture. The ultimate goal is to design algorithms so that $\lim_{T\to\infty} \operatorname{regret}(T)/T = 0$, i.e., the average regret over the rounds converges to 0. Such an algorithm is said to have sublinear regret as $\operatorname{regret}(T) = o(T)$.

Lemma 1. *Consider a convex function $f : S \to \mathbb{R}$ and assume that f has a subgradient oracle and minimizer $x_* = \operatorname{argmin}_{x\in S} f(x)$. If $p_t \in \partial f(x_t)$ for every $t \geq 0$, letting $f_* = f(x_*)$, then we have*

$$f\left(\frac{1}{T}\sum_{t=0}^{T-1} x_t\right) - f_* \leq \frac{\operatorname{regret}(T)}{T}.$$

Proof. It follows from the defition of regret in (1) and the convexity of f that

$$\begin{aligned} f\left(\frac{1}{T}\sum_{t=0}^{T-1} x_t\right) - f_* &\leq \frac{1}{T}\sum_{t=0}^{T-1} f(x_t) - f_* \leq \frac{1}{T}\sum_{t=0}^{T-1} p_t^\top (x_t - x_*) \\ &\leq \frac{1}{T}\left[\sum_{t=0}^{T-1} p_t^\top x_t - \min_{x\in S}\sum_{t=0}^{T-1} p_t^\top x\right] = \frac{\operatorname{regret}(T)}{T}. \end{aligned}$$

□

We note that the above lemma gives an important reduction: if we have an algorithm to solve online linear optimization, given $p_t \in \partial f(x_t)$ for every $t \geq 0$, then the same algorithm also solves the offline optimization $\min_{x \in S} f(x)$.

2 Follow the regularized leader

In online linear optimization, a standard algorithm is the follow the regularized leader (FTRL): pick a regularizer $r : \mathbb{R}^n \to \mathbb{R}$, that is differentiable and μ-strongly convex on S with respect to some norm $\|\cdot\|$, then FTRL generates x_T as follows.

Algorithm 1 FTRL

Input: pick an initial strategy $x_0 \in S$

for $T \geq 1$ **do**

Step 1. Enviroment reveals a penalty $p_{T-1} \in \mathbb{R}^n$.

Step 2. Compute $x_T = \underset{x \in S}{\operatorname{argmin}} \left\{ \Phi_T(x) := \eta \sum_{t=0}^{T-1} p_t^\top x + r(x) \right\}$.

end for

Note that Φ_T is defined as in FTRL for $T \geq 1$ and $\Phi_0(x) = r(x)$. We assume there exists an oracle to solve the subproblem in step 2 of FTRL. For example, if $r(x)$ is ℓ_2-norm squared, then the oracle for step 2 is just proj_S. It is interesting to observe that FTRL is the same as the constant stepsize dual averaging (see Lecture 9). As a result, the regret analysis of FTRL is very close to the convergence analysis of dual averaging, while we still present the analysis for completeness.

Lemma 2. *If $\|p_t\|_* \leq G$ for all $t \geq 0$, applying FTRL on the online linear optimization problem, then for every $T \geq 0$, we have*

$$-\frac{\eta^2 G^2 T}{2\mu} + \eta \sum_{t=0}^{T-1} p_t^\top x_t + \min_{x \in S} r(x) \leq \min_{x \in S} \Phi_T(x).$$

Proof. Proof by induction. Since $\Phi_T(x) = r(x)$, the case $T = 0$ is trivial. Assume the claim is true for some $T \geq 0$. Using the definition of Φ_T in FTRL and the induction hypothesis, we have

$$\begin{aligned}
\Phi_{T+1}(x_{T+1}) &= \Phi_T(x_{T+1}) + \eta p_T^\top x_{T+1} \\
&\geq \Phi_T(x_T) + \frac{\mu}{2}\|x_{T+1} - x_T\|^2 + \eta p_T^\top x_{T+1} \\
&\geq -\frac{\eta^2 G^2 T}{2\mu} + \eta \sum_{t=0}^{T-1} p_t^\top x_t + \min_{x \in S} r(x) + \eta \left[p_T^\top x_{T+1} + \frac{\mu}{2\eta}\|x_{T+1} - x_T\|^2 \right],
\end{aligned}$$

where the first inequality is due to the fact that Φ_T is μ-strongly convex in $\|\cdot\|$. It follows from the Cauchy-Schwarz inequality and the assumption that $\|p_t\|_* \leq G$ that

$$p_T^\top x_{T+1} + \frac{\mu}{2\eta}\|x_{T+1} - x_T\|^2 \geq -\|p_T\|_* \|x_{T+1} - x_T\| + \frac{\mu}{2\eta}\|x_{T+1} - x_T\|^2 + p_T^\top x_T \geq -\frac{\eta G^2}{2\mu} + p_T^\top x_T.$$

Therefore, combining the above two equations, we verify that the claim holds for the case $T+1$. $\square$

Theorem 1. *If $\|p_t\|_* \le G$ for every $t \ge 0$, applying FTRL on the online linear optimization problem, then for every $T \ge 1$, we have the following regret bound*

$$\mathrm{regret}(T) \le \frac{\eta G^2 T}{2\mu} + \frac{1}{\eta}\left(\max_{x\in S} r(x) - \min_{x\in S} r(x)\right).$$

Moreover, if $p_t \in \partial f(x_t)$ for every $t \ge 0$ and $\max_{x\in S} r(x) - \min_{x\in S} r(x) \le R^2$ for some $R > 0$, and using stepsize $\eta = \sqrt{(2R^2\mu)/(TG^2)}$, then the regret of FTRL is bounded as follows

$$\mathrm{regret}(T) \le RG\sqrt{\frac{2T}{\mu}},$$

and FTRL gives an optimality gap of $\min_{x\in S} f(x)$ as follows

$$f\left(\frac{1}{T}\sum_{t=0}^{T-1} x_t\right) - f_* \le RG\sqrt{\frac{2}{T\mu}}.$$

Proof. It follows from Lemma 2 and the definition of Φ_T that for every $x \in S$

$$-\frac{\eta^2 G^2 T}{2\mu} + \eta\sum_{t=0}^{T-1} p_t^\top x_t + \min_{x\in S} r(x) \le \eta\sum_{t=0}^{T-1} p_t^\top x + r(x).$$

Rearranging the terms, we have

$$\sum_{t=0}^{T-1} p_t^\top (x_t - x) \le \frac{\eta G^2 T}{2\mu} + \frac{1}{\eta}\left(r(x) - \min_{x\in S} r(x)\right).$$

Maximizing over $x \in S$ and using the definition of regret in (1), we obtain

$$\mathrm{regret}(T) \le \frac{\eta G^2 T}{2\mu} + \frac{1}{\eta}\left(\max_{x\in S} r(x) - \min_{x\in S} r(x)\right).$$

Using stepsize $\eta = \sqrt{(2R^2\mu)/(TG^2)}$ and bound R^2, we derive the regret of FTRL bounded by

$$\mathrm{regret}(T) \le RG\sqrt{\frac{2T}{\mu}}.$$

It immediately follows from the above regret of FTRL and Lemma 1 that

$$f\left(\frac{1}{T}\sum_{t=0}^{T-1} x_t\right) - f_* \le \frac{\mathrm{regret}(T)}{T} = \frac{RG\sqrt{2T/\mu}}{T} = RG\sqrt{\frac{2}{T\mu}}.$$

$\square$

Clearly, the complexity for FTRL to find an ε-solution to $\min_{x \in S} f(x)$ is

$$\mathcal{O}\left(\frac{R^2 G^2}{\mu \varepsilon^2}\right),$$

which matches that of dual averaging for solving nonsmooth optimization problems.

As we know dual averaging is also applicable to smooth optimization problems, we now use FTRL to recover the convergence of projected gradient-type methods (such as dual averaging and the projected gradient method) for minimizing smooth functions.

Lemma 3. *Assume that f is L-smooth over S, and let $r(x) = \|x - x_0\|_2^2/2$ in step 2 of FTRL and let $p_t = \nabla f(x_t)$ for every $t \geq 0$. Applying FTRL with $\eta = 1/L$ on the online linear optimization problem, then for every $T \geq 0$, we have*

$$\eta \sum_{t=1}^{T} f(x_t) - \eta \sum_{t=0}^{T-1} \left(f(x_t) - p_t^\top x_t \right) \leq \min_{x \in S} \Phi_T(x)$$

Proof. Proof by induction. Since $\Phi_0(x) = \|x - x_0\|_2^2/2$, the case $T = 0$ is trivial. Assume the claim is true for some $T \geq 0$. Using the definition of Φ_T in FTRL and the induction hypothesis, we have

$$\begin{aligned}
\Phi_{T+1}(x_{T+1}) &= \Phi_T(x_{T+1}) + \eta p_T^\top x_{T+1} \\
&\geq \Phi_T(x_T) + \frac{1}{2}\|x_{T+1} - x_T\|_2^2 + \eta p_T^\top x_{T+1} \\
&\geq \eta \sum_{t=1}^{T} f(x_t) - \eta \sum_{t=0}^{T-1} \left(f(x_t) - p_t^\top x_t \right) + \eta \left[p_T^\top x_{T+1} + \frac{1}{2\eta}\|x_{T+1} - x_T\|_2^2 \right],
\end{aligned}$$

where the first inequality is due to the fact that Φ_T is 1-strongly convex in $\|\cdot\|_2$. It follows from the fact thats $p_t = \nabla f(x_t)$ and $\eta = 1/L$, and the assumption that f is L-smooth that

$$p_T^\top x_{T+1} + \frac{1}{2\eta}\|x_{T+1} - x_T\|_2^2 = \nabla f(x_T)^\top x_{T+1} + \frac{L}{2}\|x_{T+1} - x_T\|_2^2 \geq f(x_{T+1}) - \left(f(x_T) - \nabla f(x_T)^\top x_T \right).$$

Therefore, combining the above two equations, we verify that the claim for the case $T+1$ holds. □

Theorem 2. *Assuming the conditions in Lemma 3 hold and S has a diameter $D > 0$, and applying FTRL with $\eta = 1/L$ on the online linear optimization problem, then for every $T \geq 1$, we have the following regret bound*

$$\text{regret}(T) \leq f(x_0) - f_* + \frac{LD^2}{2},$$

and FTRL gives an optimality gap of $\min_{x \in S} f(x)$ as follows

$$f\left(\frac{1}{T}\sum_{t=0}^{T-1} x_t\right) - f_* \leq \frac{2[f(x_0) - f_*] + LD^2}{2T}.$$

Proof. It follows from Lemma 3 that for every $x \in S$

$$\eta[f(x_T) - f(x_0)] + \eta \sum_{t=0}^{T-1} p_t^\top x_t \le \eta \sum_{t=0}^{T-1} p_t^\top x + \frac{1}{2}\|x - x_0\|_2^2.$$

Rearranging the terms and using the fact that $\eta = 1/L$ yields

$$\sum_{t=0}^{T-1} p_t^\top (x_t - x) \le f(x_0) - f(x_T) + \frac{L}{2}\|x - x_0\|_2^2.$$

Maximizing over $x \in S$ and using the definition of regret in (1) and the boundedness of S, we obtain

$$\text{regret}(T) \le f(x_0) - f_* + \frac{LD^2}{2}.$$

It immediately follows from the above regret of FTRL and Lemma 1 that

$$f\left(\frac{1}{T}\sum_{t=0}^{T-1} x_t\right) - f_* \le \frac{\text{regret}(T)}{T} \le \frac{2[f(x_0) - f_*] + LD^2}{2T}.$$

□

The following theorem gives a different convergence result of FTRL without using the reduction, i.e., Lemma 1. It recovers the convergence rate of dual averaging for solving smooth optimization problems.

Theorem 3. *Assuming the conditions in Lemma 3 hold and applying FTRL with $\eta = 1/L$ on the online linear optimization problem, then for every $T \ge 1$, we have*

$$f\left(\frac{1}{T}\sum_{t=1}^{T} x_t\right) - f_* \le \frac{L\|x_0 - x_*\|_2^2}{2T}.$$

Proof. It immediately follows from the fact that $p_t = \nabla f(x_t)$ for every $t \ge 0$ and Lemma 3 that for every $x \in S$

$$\eta \sum_{t=1}^{T} f(x_t) - \eta \sum_{t=0}^{T-1} \left(f(x_t) - \nabla f(x_t)^\top x_t\right) \le \eta \sum_{t=0}^{T-1} \nabla f(x_t)^\top x + \frac{1}{2}\|x - x_0\|_2^2.$$

Rearranging the terms and using the fact that $\eta = 1/L$ yields

$$\sum_{t=1}^{T} f(x_t) \le \sum_{t=0}^{T-1} (f(x_t) + \langle \nabla f(x_t), x - x_t \rangle) + \frac{L}{2}\|x - x_0\|_2^2 \le Tf(x) + \frac{L}{2}\|x - x_0\|_2^2,$$

where the second inequality is due to the convexity of f. Taking $x = x_*$ in the above inequality and using the convexity of f again, we obtain

$$f\left(\frac{1}{T}\sum_{t=1}^{T} x_t\right) - f_* \leq \frac{L\|x_0 - x_*\|_2^2}{2T}.$$

$\square$

Note that Theorem 3 shows convergence at $\frac{1}{T}\sum_{t=1}^{T} x_t$, which is different from $\frac{1}{T}\sum_{t=0}^{T-1} x_t$ in Theorems 1 and 2.

3 Saddle point problem

In Section 2, we have seen FTRL originally designed for online linear optimization can be readily applied to offline optimization, both smooth and nonsmooth minimization problems. We discuss another interesting application of FTRL in this section, that is, the saddle point problem.

The saddle point problem considered in this section is as follows

$$\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} f(x, y)$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \in \mathbb{R}^m$ are closed convex sets and $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ is differentiable and convex-concave (i.e., $f(\cdot, y)$ is convex for any $y \in \mathcal{Y}$ and $f(x, \cdot)$ is concave for any $x \in \mathcal{X}$). It is equivalent to

$$\min_{x\in\mathcal{X}} \left\{ f_{\mathcal{X}}(x) := \max_{y\in\mathcal{Y}} f(x, y) \right\}.$$

For simplicity, we denote $f(z) = f(z_{\mathcal{X}}, z_{\mathcal{Y}})$, where $z \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ and $z_{\mathcal{X}}$ (resp., $z_{\mathcal{Y}}$) denotes the $\mathcal{X}$-component (resp., $\mathcal{Y}$-component) of z.

For every $z \in \mathcal{Z}$, we define the duality gap as

$$\mathrm{gap}(z) = \mathrm{gap}(z_{\mathcal{X}}, z_{\mathcal{Y}}) = \max_{y\in Y} f(z_{\mathcal{X}}, y) - \min_{x\in\mathcal{X}} f(x, z_{\mathcal{Y}}) = f_{\mathcal{X}}\left(z_{\mathcal{X}}\right) - f_{\mathcal{Y}}\left(z_{\mathcal{Y}}\right),$$

where $f_{\mathcal{Y}}(y) := \min_{x\in\mathcal{X}} f(x, y)$. We say that $z \in \mathcal{Z}$ is an ε-Nash equilibrium, or has ε-duality gap, or is ε-optimal if $\mathrm{gap}(z) \leq \varepsilon$. Further, we call $z \in \mathcal{Z}$ a Nash equilibrium if $\mathrm{gap}(z) = 0$. An intuitive interpretation of the duality gap is as follows: consider there are two players x and y in a zero-sum game, the goal of x is to minimize $f(x, y)$ and hence save cost, while the goal of y is to maximize $f(x, y)$ and hence make profit. The duality gap can be written as

$$\mathrm{gap}(z) = \left[\max_{y\in Y} f\left(z_{\mathcal{X}}, y\right) - f(z)\right] + \left[f(z) - \min_{x\in\mathcal{X}} f\left(x, z_{\mathcal{Y}}\right)\right],$$

where the first gap term is the profit that y makes with $z_{\mathcal{X}}$ being fixed and the second gap term is the cost that x saves with $z_{\mathcal{Y}}$ being fixed. A Nash equilibrium is a pair of stategies z at which both x and y are not willing to change their individual strategies $z_{\mathcal{X}}$ and $z_{\mathcal{Y}}$, as none of them would increase their own utility functions (i.e., saving cost and making profit) by doing so.

Note that $f_{\mathcal{X}}(z_{\mathcal{X}}) \geq f(z) \geq f_{\mathcal{Y}}(z_{\mathcal{Y}})$ for every $z \in \mathcal{Z}$ by definition, and hence $f_{\mathcal{X}}^* \geq f_{\mathcal{Y}}^*$ where $f_{\mathcal{X}}^* = \min_{x\in\mathcal{X}} f_{\mathcal{X}}(x)$ and $f_{\mathcal{Y}}^* = \max_{y\in\mathcal{Y}} f_{\mathcal{Y}}(y)$. This is indeed the weak duality and it is equivalent to

$$\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} f(x,y) \geq \max_{y\in\mathcal{Y}} \min_{x\in\mathcal{X}} f(x,y).$$

If $f(x,y)$ is convex-concave, then "=" holds. Clearly, $\mathrm{gap}(z) \geq f_{\mathcal{X}}(z_{\mathcal{X}}) - f_{\mathcal{X}}^* + f_{\mathcal{Y}}^* - f_{\mathcal{Y}}(z_{\mathcal{Y}})$ and therefore if $z \in \mathcal{Z}$ is an ε-Nash equilibrium, then $z_{\mathcal{X}}$ (resp., $z_{\mathcal{Y}}$) is an ε-solution for $\min_{x\in\mathcal{X}} f_{\mathcal{X}}$ (resp., $\max_{y\in\mathcal{Y}} f_{\mathcal{Y}}$). Furthermore, this reasoning implies that if there exists a Nash equilibrium, i.e., $z^* \in \mathcal{Z}$ with $\mathrm{gap}\,(z^*) = 0$, then $f_{\mathcal{X}}\left(z_{\mathcal{X}}^*\right) = f_{\mathcal{X}}^*$ and $f_{\mathcal{Y}}\left(z_{\mathcal{Y}}^*\right) = f_{\mathcal{Y}}^*$ and correspondingly $f_{\mathcal{X}}^* = f_{\mathcal{Y}}^*$. This is called the strong duality.

The discussions above reflect some key ideas in Lecture 9.

Analogous to the reduction from online linear optimization to offline optimization (see Lemma 1), we next develop the reduction from online linear optimization to the saddle point problem. Thus, we can apply FTRL to solve the saddle point problem.

Lemma 4. *Let $g : \mathcal{X}\times\mathcal{Y} \to \mathcal{X}\times\mathcal{Y}$ denote the gradient for $f(x,y)$, i.e., $g(x,y) = (\nabla_{\mathcal{X}} f(x,y), -\nabla_{\mathcal{Y}} f(x,y))$. For $z^0, \ldots, z^{T-1} \in \mathcal{Z}$ and $\bar{z} = \frac{1}{T}\sum_{t=0}^{T-1} z^t$, we have*

$$\mathrm{gap}(\bar{z}) \leq \frac{\mathrm{regret}(T)}{T}.$$

Proof. Since f is convex-concave, we have for every $u \in \mathcal{Z}$,

$$f(u_{\mathcal{X}}, z_{\mathcal{Y}}^t) \geq f(z^t) + \nabla_{\mathcal{X}} f(z^t)^\top (u_{\mathcal{X}} - z_X^t),$$
$$f(z_{\mathcal{X}}^t, u_{\mathcal{Y}}) \leq f(z^t) + \nabla_{\mathcal{Y}} f(z^t)^\top (u_{\mathcal{Y}} - z_Y^t).$$

Combining the above two inequalities yields that

$$g(z^t)^\top (z^t - u) \geq f(z_{\mathcal{X}}^t, u_{\mathcal{Y}}) - f(u_{\mathcal{X}}, z_{\mathcal{Y}}^t).$$

Summing the above inequality from $t = 0$ to $T-1$ and applying Jensen's inequality, we have

$$\frac{1}{T}\sum_{t=0}^{T-1} g(z^t)^\top (z^t - u) \geq \frac{1}{T}\sum_{t=0}^{T-1} [f(z_{\mathcal{X}}^t, u_{\mathcal{Y}}) - f(u_{\mathcal{X}}, z_{\mathcal{Y}}^t)] \geq f(\bar{z}_{\mathcal{X}}, u_{\mathcal{Y}}) - f(u_{\mathcal{X}}, \bar{z}_{\mathcal{Y}}).$$

Maximizing over $u \in \mathcal{Z}$ and using the definition of regret in (1), we obtain

$$\frac{\mathrm{regret}(T)}{T} \geq \max_{u\in\mathcal{Z}} f(\bar{z}_{\mathcal{X}}, u_{\mathcal{Y}}) - f(u_{\mathcal{X}}, \bar{z}_{\mathcal{Y}}).$$

It follows from the definition of the duality gap that

$$\mathrm{gap}(\bar{z}) = \max_{y\in\mathcal{Y}} f(\bar{z}_{\mathcal{X}}, y) - \min_{x\in\mathcal{X}} f(x, \bar{z}_{\mathcal{Y}}) = \max_{u\in\mathcal{Z}} f(\bar{z}_{\mathcal{X}}, u_{\mathcal{Y}}) - f(u_{\mathcal{X}}, \bar{z}_{\mathcal{Y}}).$$

The conclusion of the lemma immediately follows. □

We skip the regret of FTRL and its implication in the saddle point problem for shortness.

4 Online mirror descent

As discussed in Lecture 9, dual averaging is quite similar to mirror descent. Interestingly, mirror descent also has an online counterpart analogous to FTRL, which is known as online mirror descent (OMD).

Consider the online linear optimization problem and OMD is as follows.

Algorithm 2 OMD

Input: pick an initial strategy $x_0 \in S$, stepsize $\eta > 0$.
for $t \geq 0$ **do**
  Step 1. Enviroment reveals a penalty $p_t \in \mathbb{R}^n$.
  Step 2. Compute $x_{t+1} = \operatorname*{argmin}_{x \in S} \left\{ \eta p_t^\top x + D_w(x, x_t) \right\}$.
end for

Theorem 4. *Assume w is differentiable and ρ-strongly convex w.r.t. $\|\cdot\|$ and $\|p_t\|_* \leq G$ for every $t \geq 0$. Derive the regret bound of OMD*

$$\text{regret}(T) \leq \frac{\max_{x \in S} D_w(x, x_0)}{\eta} + \frac{\eta G^2 T}{2\rho}.$$

Further, assume $\max_{x \in S} D_w(x, x_0) \leq R^2$ for some $R > 0$ and take $\eta = \sqrt{2\rho R^2}/\sqrt{G^2 T}$, then show the regret becomes

$$\text{regret}(T) \leq \frac{\sqrt{2T} RG}{\sqrt{\rho}}.$$

The proof is left as a homework problem. In general, for the online convex optimization problem, we replace p_t in OMD and FTRL by $\nabla f_t(x_t)$, and can prove similar regret bounds in terms of the general regret in (2). We next present an interesting application of the (generalized)OMD in proving the minimax theorem.

Theorem 5. *Let $\mathcal{X} \in \mathbb{R}^n$ and $\mathcal{Y} \in \mathbb{R}^m$ be compact convex sets. Let $f(x, y)$ be continuous and convex-concave, with some upper bound G on the partial subgradients with respect to x and y. Then, we have*

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y) = \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} f(x, y).$$

Proof. We have shown in Section 3 that

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y) \geq \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} f(x, y).$$

We next show the other direction by the regret of OMD. We run a repeated game where the players choose a strategy x_t, y_t at each round t. The x player chooses x_t according to OMD, while y_t is always chosen as

$$y_t = \operatorname{argmax}_{y \in \mathcal{Y}} f(x_t, y). \tag{3}$$

Let the average strategies be

$$\bar{x} = \frac{1}{T}\sum_{t=0}^{T-1} x_t, \quad \bar{y} = \frac{1}{T}\sum_{t=0}^{T-1} y_t.$$

For the x player, using a general version of Theorem 4, we have

$$\text{regret}(T) = \sum_{t=0}^{T-1} f(x_t, y_t) - \min_{x\in\mathcal{X}} \sum_{t=0}^{T-1} f(x, y_t) \le \frac{\sqrt{2T}RG}{\sqrt{\rho}}. \tag{4}$$

Recall that $f_{\mathcal{X}}(x) = \max_{y\in\mathcal{Y}} f(x,y)$ and $f_{\mathcal{Y}}(y) = \min_{x\in\mathcal{X}} f(x,y)$. Noting that $f_{\mathcal{X}}(x)$ is convex and using the concavity of $f(x,\cdot)$, we have

$$\min_{x\in\mathcal{X}}\max_{y\in\mathcal{Y}} f(x,y) = \min_{x\in\mathcal{X}} f_{\mathcal{X}}(x) \le f_{\mathcal{X}}(\bar{x}) \le \frac{1}{T}\sum_{t=0}^{T-1} f_{\mathcal{X}}(x_t), \tag{5}$$

$$\max_{y\in\mathcal{Y}}\min_{x\in\mathcal{X}} f(x,y) = \max_{y\in\mathcal{Y}} f_{\mathcal{Y}}(y) \ge f_{\mathcal{Y}}(\bar{y}) = \min_{x\in\mathcal{X}} f(x,\bar{y}) \ge \min_{x\in\mathcal{X}} \frac{1}{T}\sum_{t=0}^{T-1} f(x, y_t). \tag{6}$$

Moreover, it follows from (3) that

$$f(x_t, y_t) = \max_{y\in\mathcal{Y}} f(x_t, y) = f_{\mathcal{X}}(x_t).$$

Combining the above relation and (5), we obtain

$$\min_{x\in\mathcal{X}}\max_{y\in\mathcal{Y}} f(x,y) \le \frac{1}{T}\sum_{t=0}^{T-1} f(x_t, y_t).$$

Plugging this inequality and (6) into the regret (4), we have

$$\begin{aligned}\min_{x\in\mathcal{X}}\max_{y\in\mathcal{Y}} f(x,y) - \max_{y\in\mathcal{Y}}\min_{x\in\mathcal{X}} f(x,y) &\le \frac{1}{T}\sum_{t=0}^{T-1} f(x_t, y_t) - \frac{1}{T}\min_{x\in\mathcal{X}}\sum_{t=0}^{T-1} f(x, y_t) \\ &= \frac{\text{regret}(T)}{T} \overset{(4)}{\le} \frac{\sqrt{2}RG}{\sqrt{\rho T}}.\end{aligned}$$

Taking the limit $T \to \infty$, we have

$$\min_{x\in\mathcal{X}}\max_{y\in\mathcal{Y}} f(x,y) \le \max_{y\in\mathcal{Y}}\min_{x\in\mathcal{X}} f(x,y).$$

Therefore, "=" holds and the theorem is proved. $\square$

5 Fisher market equilibrium

We conclude this lecture with an interesting application in economics, where standard mirror descent fails, but succeeds under the notion of relative smoothness.

We consider the economy of a Fisher market, where there are n buyers, each with some budget $B_i > 0$ (fake currency), and m items to be allocated. Without loss of generality, we assume that each item is of unit supply, one. In the context of fair division and market equilibrium, we wish to find a set of prices $p \in \mathbb{R}_+^m$ for each of the m items and an allocation $x \in \mathbb{R}_+^{n\times m}$ for each of the nm buyer-item pair such that the market clears. Intuitively, a market clears when there exists an allocation x such that everybody is assigned an optimal allocation given the prices and their budget, and all items are exactly allocated at their supply. "Optimal allocation" is defined so as to maximize the utility functions of the buyers, which is assumed to be linear $u_i(x_i) = \langle v_i, x_i \rangle$ for $i = 1, \ldots, n$, where $v_i \in \mathbb{R}^m$ is buyer i's unit utility/value of all items. We assume the items are infinitely-divisible among buyers. For brevity, we omit the definition of fair division in this lecture and directly present the Eisenberg-Gale (EG) convex program that charactorizes this market equilibrium:

$$\begin{aligned} \max_{x \in \mathbb{R}_+^{m\times n}} \quad & \textstyle\sum_{i\in[n]} B_i \log \langle v_i, x_i \rangle \\ \text{s.t.} \quad & \textstyle\sum_{i\in[n]} x_{ij} \le 1, \quad \forall j = 1, \ldots, m, \end{aligned} \tag{EG}$$

where x_{ij} is the (percentage) allocation of item j to buyer i, and x_i is the allocation of all items to buyer i. The associated dual variable to each constraint is the price p_j of item j. The solution to the primal problem x along with the dual variable p yields a market equilibrium.

Formally, the demand set of a buyer i with budget B_i is

$$D_i(p) = \operatorname*{argmax}_{x_i \ge 0} u_i(x_i) \ \text{ s.t. } \ \langle p, x_i \rangle \le B_i. \tag{7}$$

Definition 1. *A market equilibrium is an allocation $x \in \mathbb{R}_{\ge 0}^{n\times m}$ and a price vector $p \in \mathbb{R}_{\ge 0}^m$ such that:*

(i) Demands are satisfied: $x_i \in D(p)$ for all buyers i.

(ii) The market clears: $\sum_{i\in[n]} x_{ij} \le 1$, and $\sum_{i\in[n]} x_{ij} = 1$ if $p_j > 0$.

5.1 Duals of the Eisenberg-Gale convex program

Rewriting the EG program as

$$\begin{aligned} \max_{x \in \mathbb{R}_+^{n\times m}, u \in \mathbb{R}^n} \quad & \textstyle\sum_{i\in[n]} B_i \log u_i \\ \text{s.t.} \quad & u_i \le \langle v_i, x_i \rangle, \qquad \forall i = 1, \ldots, n, \\ & \textstyle\sum_{i\in[n]} x_{ij} \le 1, \quad \forall j = 1, \ldots, m. \end{aligned}$$

Introducing dual variables β_i (corresponding to the utility price of buyer i), and p_j (the price of item j), the Lagrangian is

$$L(x, u, \beta, p) = \sum_{i \in [n]} B_i \log u_i + \sum_{i \in [n]} \beta_i \left(\langle v_i, x_i \rangle - u_i \right) + \sum_{j \in [m]} p_j \left(1 - \sum_{i \in [n]} x_{ij} \right).$$

The standard Lagrangian dual is then

$$\min_{p \in \mathbb{R}_+^m, \beta \in \mathbb{R}_+^n} \max_{x \in \mathbb{R}_+^{n \times m}, u} L(x, u, \beta, p).$$

Maximizing over x and u, the primal objective function of (p, β) becomes

$$\begin{aligned}
\max_{x \in \mathbb{R}_+^{n \times m}, u} L(x, u, \beta, p) &= \sum_{j \in [m]} p_j + \sum_i \left[\max_{u_i} \left(B_i \log u_i - \beta_i u_i \right) + \max_{x_i \geq 0} \langle \beta_i v_i - p, x_i \rangle \right] \\
&= \sum_{j \in [m]} p_j + \sum_i \left[\max_{u_i} \left(B_i \log u_i - \beta_i u_i \right) + I_{\beta_i v_i \leq p} \right] \\
&= \sum_{j \in [m]} p_j + \sum_i \left[B_i \max_{u_i} \left(\log u_i - \frac{\beta_i}{B_i} u_i \right) + I_{\beta_i v_i \leq p} \right] \\
&= \sum_{j \in [m]} p_j + \sum_i \left[B_i \left(-1 - \log \beta_i + \log B_i \right) + I_{\beta_i v_i \leq p} \right].
\end{aligned}$$

Thus, dropping the constant term $\sum_{i \in [n]} (\log B_i - B_i)$, we obtain the dual of EG

$$\begin{array}{lll}
\min\limits_{p \in \mathbb{R}_+^m, \beta \in \mathbb{R}^n} & \sum_j p_j - \sum_{i \in [n]} B_i \log \beta_i & \\
\text{s.t.} & p_j \geq v_{ij} \beta_i, & \forall i, j.
\end{array} \tag{dual EG}$$

The EG dual can be used to derive several attractive algorithms for computing Fisher market equilibria. Many of these algorithms are based on the fact that we can reformulate away either β or p in (dual EG).

For a fixed set of prices p, we take β_i as large as possible $\beta_i = p_j / v_{ij}$ and hence obtain the program in p

$$\min_{p \in \mathbb{R}_+^m} \sum_j p_j - \sum_{i \in [n]} B_i \log \left(\min_{j \in [m]} \frac{p_j}{v_{ij}} \right).$$

Noting the above objective function is convex in p and amipplying the subgradient method, we get the following price update dynamics:

$$p^{k+1} = p^k + \eta \left(\sum_{i \in [n]} D_i(p^k) - \mathbf{1} \right),$$

where $\mathbf{1}$ is an all-one vector, and D_i is as in (7) and should be understood as selecting an arbitrary element of the demand set for each buyer i. This derivation is left as a homework problem. This dynamics gives an intuitive economics explanation. Noting that $\sum_{i\in[n]} D_i(p_j^k)$ is the demand for item j and 1 is the supply for item j, so if item j is overdemanded, then the market increases price p_j, and vice versa.

For a fixed β, each price p_j should be made as small as possible. This yields the following program in β

$$\min_{\beta\in\mathbb{R}_+^n} \sum_{j\in[m]} \max_{i\in[n]} \beta_i v_{ij} - \sum_{i\in[n]} B_i \log(\beta_i).$$

This program can be used to derive a very natural (first-price) auction-based dynamics, known as the PACE dynamics, which arises by applying the follow-the-leader (FTL). Note this is related to but different from FTRL introduced in early sections.

5.2 Shmyrev's convex program and proportional response dynamics

Introducing a change of variables to (dual EG), by letting $q_j = \log p_j$ and $\gamma_i = -\log \beta_i$, we get

$$\begin{array}{ll} \min\limits_{q,\gamma} & \sum_j e^{q_j} + \sum_{i\in[n]} B_i\gamma_i \\ \text{s.t.} & q_j + \gamma_i \geq \log v_{ij}, \qquad \forall i, j. \end{array}$$

Now we introduce the Lagrangian multiplier $b_{ij} \geq 0$ for the constraint above to get the following dual

$$\begin{aligned} &\max_{b\in\mathbb{R}_+^{n\times m}} \min_{q,\gamma} \sum_j e^{q_j} + \sum_{i\in[n]} B_i\gamma_i + \sum_{ij} b_{ij}(\log v_{ij} - q_j - \gamma_i) \\ =& \max_{b\in\mathbb{R}_+^{n\times m}} \left[\sum_{ij} b_{ij}\log v_{ij} + \sum_{j\in[m]} \min_{q_j}\left[e^{q_j} - \sum_{i\in[n]} b_{ij}q_j\right] + \sum_{i\in[n]} \min_{\gamma_i} \gamma_i\left[B_i - \sum_{j\in[m]} b_{ij}\right]\right]. \end{aligned}$$

Noting that the inner minimization enforces

$$p_j = e^{q_j} = \sum_{i\in[n]} b_{ij}, \quad B_i = \sum_{j\in[m]} b_{ij},$$

we arrive at the Shmyrev's convex program

$$\begin{array}{ll} \max\limits_{b\in\mathbb{R}_+^{n\times m}} & \sum_{ij} b_{ij}\log v_{ij} + \sum_{j\in[m]} (p_j - p_j\log p_j) \\ \text{s.t.} & \sum_{i\in[n]} b_{ij} = p_j, \quad j = 1,\ldots,m, \\ & \sum_{j\in[m]} b_{ij} = B_i, \quad i = 1,\ldots,n. \end{array}$$

The variable b_{ij} can be interpreted as how much of their budget buyer i spends on a given item j. Since $\sum_{j\in[m]} p_j = \sum_{i\in[n]} B_i$ is a constant it does not affect the objective, so we rewrite Shmyrev's convex program as

$$
\begin{array}{lll}
\max\limits_{b \in \mathbb{R}_+^{n\times m}} & \sum_{ij} b_{ij} \log v_{ij} - \sum_{j\in[m]} p_j \log p_j & \\
\text{s.t.} & \sum_{i\in[n]} b_{ij} = p_j, \quad j = 1, \ldots, m, & \\
& \sum_{j\in[m]} b_{ij} = B_i, \quad i = 1, \ldots, n. &
\end{array} \tag{Shmyrev}
$$

Denoting $p_j(b) = \sum_{i\in[n]} b_{ij}$, we rewrite (Shmyrev) as a constrained optimization $\min\{f(b) : b \in Q\}$ in terms of b_{ij} only, where

$$
f(b) = -\sum_{ij} b_{ij} \log v_{ij} + \sum_{j\in[m]} p_j(b) \log p_j(b) = -\sum_{ij} b_{ij} \log\left(v_{ij}/p_j(b)\right), \tag{8}
$$

and

$$
Q = \left\{ b \in \mathbb{R}_+^{n\times m} : \sum_{j\in[m]} b_{ij} = B_i, \quad i = 1, \ldots, n \right\}.
$$

Now, noting that $\nabla_{ij} f(b) = 1 - \log(v_{ij}/p_j(b))$, and choosing the distance generating function as the negative entropy $w(b) = \sum_{ij} b_{ij} \log b_{ij}$ and stepsize $\eta = 1$, we have the mirror descent update (see PS3 P4(d) for the multiplicative update)

$$
b_{ij}^{k+1} \propto b_{ij}^k \exp\left(-1 + \log\left(v_{ij}/p_j(b)\right)\right) \propto b_{ij}^k \left(v_{ij}/p_j(b)\right) = \frac{1}{Z} b_{ij}^k \left(v_{ij}/p_j(b)\right),
$$

where Z is a normalization constant such that $\sum_{j\in[m]} b_{ij}^{k+1} = B_i$.

Amazingly, mirror descent on (Shmyrev) with stepsize $\eta = 1$ becomes the following very natural algorithm, namely the proportional response dynamics:

- at each iteration k, each buyer i submits a bid vector b_i^k (recommended by MD update);
- given the bids, a price $p_j^k = p_j(b^k) = \sum_{i\in[n]} b_{ij}^k$ is computed for each item j;
- each buyer is allocated a fraction
$$
x_{ij}^k = \frac{b_{ij}^k}{p_j^k}
$$
of each item (this is intuitive, but can also be seen from KKT conditions of (EG));
- each buyer then submits their next bid on item j proportional to the utility received from item j in iteration k:
$$
b_{ij}^{k+1} = B_i \frac{x_{ij}^k v_{ij}}{\sum_{j'} x_{ij'}^k v_{ij'}}.
$$

Definition 2. *We say f is L-smooth relative to h on Q if for any $x, y \in \operatorname{int} Q$, there is a scalar L for which*

$$
f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + L D_h(y, x),
$$

where D_h is the Bregman divergence of h.

In view of this notion of relative smoothness, it can be shown that $f(b)$ as in (8) is 1-smooth relative to the negative entropy $w(b)$.

Theorem 6. *Applying mirror desent on* (Shmyrev) *with distance generating function $w(b)$, stepsize $\eta = 1$, and b^0 being $b_{ij} = B_i/m$, then we have*

$$f(b^k) - f(b_*) \leq \frac{1}{k} \sum_i B_i \log \frac{B_i}{m}.$$

Note that by recognizing the relative smoothness of f, we essentially establish the standard $\mathcal{O}(1/k)$ convergence rate for convex smooth optimization.