Two-Player Zero-Sum Games, Minimax Theorem, Strong LP Duality.

Def (Two-Player Zero-Sum Game):

A zero-sum game with payoff matrix $M \in \mathbb{R}^{m \times n}$ is the following game:

Two players: row player and column player.

Row player plays action $i \in [m]$

Column player plays action $j \in [n]$

Row player gains $M_{i,j}$, column player loses $M_{i,j}$.

(Column player pays $M_{i,j}$ to row player.)

Example (Rock Paper Scissors):

$$M^{3 \times 3} = \begin{array}{c|ccc|} & R & P & S \\ \hline R & 0 & -1 & +1 \\ P & +1 & 0 & -1 \\ S & -1 & +1 & 0 \\ \hline \end{array}$$

Suppose column player is the adversary and knows your strategy.

Then your strategy should be randomized.

In RPS case, uniform between R/P/S: $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. Expected payoff $\emptyset$.

Row player strategy $\vec{p} \in \Delta_m$, column player strategy $\vec{q} \in \Delta_n$

Then expected payoff is $\sum_{i,j} p_i q_j M_{ij} = \vec{p}^T M \vec{q}$

Given a row strategy $\vec{p}$, column player should play $\vec{q}$ to minimize $\vec{p}^T M \vec{q}$.

Should only play entry that minimizes $\vec{p}^T M$.

$$C(p) = \min_{q \in \Delta_n} p^T M q = \min_{j \in [n]} p^T M e_j.$$

Row player wants to play $\vec{p}$ that maximizes $C(p)$ (best column player response).

$q\in\Delta_n$

Row player wants to play $\vec{p}$ that maximizes $C(p)$ (best column player response).
So row player's optimum is $\max\limits_{p\in\Delta_m} C(p)$.

Similarly, given a column strategy $\vec{q}$, row player should maximize $\vec{p}^T M \vec{q}$.

$$R(q) = \max_{p\in\Delta_m} p^T M q = \max_{i\in[m]} e_i^T M$$

Column player wants to play $\vec{q}$ that minimizes $R(p)$, optimum $\min\limits_{q\in\Delta_n} R(q)$.

Lemma (Weak Minimax Theorem): For any $p\in\Delta_m$, $q\in\Delta_n$, $C(p)\le R(q)$.

Proof: Intuitively, $R(q)$ fixes column strategy first, giving more power to row player.

Formally,
$$p^T M q \ge \min_{q'} p^T M q' = C(p),$$
$$p^T M q \le \max_{p'} p'^T M q = R(q).$$

Theorem (Von Neumann's Minimax Theorem): For any payoff matrix $M\in\mathbb{R}^{m\times n}$,

$$\max_{p\in\Delta_m} C(p) = \min_{q\in\Delta_n} R(q), \text{ i.e. } \max_{p\in\Delta_m}\min_{q\in\Delta_n} p^T M q = \min_{q\in\Delta_n}\max_{p\in\Delta_m} p^T M q$$

Doesn't matter who plays first if both play optimally!

Proof: For any $\varepsilon>0$, we establish strategies $\bar{p}, \bar{q}$ s.t. $C(\bar{p})\ge R(\bar{q})-\varepsilon$.

$$\boxed{\begin{array}{l} \bar{p} = \frac{1}{T}\sum_{t=1}^{T} p_t \text{ where } p_t \text{ is probability vector given by Hedge.} \\ \bar{q} = \frac{1}{T}\sum_{t=1}^{T} q_t \text{ where } q_t \text{ is best column response to } p_t. \end{array}}$$

Scale M down until $M\in[-1,1]^{m\times n}$.

Run Hedge with experts $[n]$. On each iteration with probability vector p_t, play gain vector Mq_t where $q_t = \operatorname{argmin}\limits_{q\in\Delta_m} p_t^T M q$.

Hedge guarantees that

$$\ldots \max \sum^{T} \langle e_i, g_t\rangle - \varepsilon T \quad \text{if } T = O\left(\frac{\ln N}{\varepsilon^2}\right).$$

Hedge guarantees that

$$\underbrace{\sum_{t=1}^{T} \langle p_t, g_t \rangle}_{\text{Hedge's gain}} \geq \underbrace{\max_{i \in [n]} \sum_{t=1}^{T} \langle e_i, g_t \rangle}_{\text{best expert response}} - \varepsilon T \quad \text{if } T = O\left(\frac{\ln N}{\varepsilon^2}\right).$$

So
$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^{T} \langle p_t, g_t \rangle &\geq \max_{i \in [n]} \frac{1}{T} \sum_{t=1}^{T} \langle e_i, g_t \rangle - \varepsilon && \text{[by Hedge]} \\
&= \max_{i \in [n]} \left\langle e_i, \frac{1}{T} \sum_{t=1}^{T} g_t \right\rangle - \varepsilon \\
&= \max_{i \in [n]} \left\langle e_i, \frac{1}{T} \sum_{t=1}^{T} M q_t \right\rangle - \varepsilon && \text{[by definition of } g_t\text{]} \\
&= \max_{i \in [n]} \left\langle e_i, M \underbrace{\left(\frac{1}{T} \sum_{t=1}^{T} q_t \right)}_{\bar{q}} \right\rangle - \varepsilon \\
&= R(\bar{q}) - \varepsilon
\end{aligned}$$

Also,
$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^{T} \langle p_t, g_t \rangle &= \frac{1}{T} \sum_{t=1}^{T} \langle p_t, M q_t \rangle && \text{[by definition of } g_t\text{]} \\
&= \frac{1}{T} \sum_{t=1}^{T} \left(\min_{q \in \Delta_m} p^T M q \right) && \text{[by choice of } q_t\text{]} \\
&\leq \min_{q \in \Delta_m} \left(\frac{1}{T} \sum_{t=1}^{T} p^T M q \right) && \text{[by concavity of min]} \\
&= \min_{q \in \Delta_m} \left(\bar{p}^T M q \right) = C(\bar{p}).
\end{aligned}$$

Define q^* best column response to $\bar{p}$. Then

$$\begin{aligned}
&\frac{1}{T}\left(\sum_{t=1}^{T} \min_{q \in \Delta_m} p^T M q \right) \\
&\leq \frac{1}{T} \sum_{t=1}^{T} p^T M q^* \\
&= \bar{p}^T M q^* = C(\bar{p}).
\end{aligned}$$

Therefore, $C(\bar{p}) \geq \frac{1}{T} \sum_{t=1}^{T} \langle p_t, g_t \rangle \geq R(\bar{q}) - \varepsilon.$

Minimax Theorem and Strong LP Duality

Minimax Theorem and Strong LP Duality

By minimax theorem, $\max\limits_{p\in\Delta_m} \min\limits_{j\in[n]} p^T M e_j = \min\limits_{q\in\Delta_n} \max\limits_{i\in[m]} e_i^T M q$

We want to convert this to LPs.

$$\max_{p\in\Delta_m} \min_{j\in[n]} p^T M e_j \iff \begin{aligned} &\max\ t \quad (= \min_{j\in[n]} e_j^T M^T p,\ \text{assume} > 0) \\ &\text{s.t. } M^T p \geq t\mathbb{1} \\ &\quad\ \ \mathbb{1}^T p = 1 \\ &\quad\ \ p \geq \vec{0} \end{aligned}$$

Convert $\mathbb{1}^T p = 1$ to $\mathbb{1}^T p \leq 1$: add new dummy p_{m+1} and row $(m+1)$ to M with all entries 0.

Now substitute $\vec{x} = \frac{1}{t}\vec{p}$:

$$\begin{aligned} &\max\ t \\ &\text{s.t. } M^T x \geq \mathbb{1} \\ &\quad\ \ \mathbb{1}^T x \leq 1/t \\ &\quad\ \ x \geq 0 \end{aligned} \iff \begin{aligned} &\min\ \mathbb{1}^T x \quad (= 1/t) \\ &\text{s.t. } M^T x \geq \mathbb{1} \\ &\quad\ \ x \geq 0. \end{aligned}$$

$$\min_{q\in\Delta_n} \max_{i\in[m]} e_i^T M q \iff \begin{aligned} &\min\ t \quad (= \max_{i\in[m]} e_i^T M q,\ \text{assume} > 0) \\ &\text{s.t. } Mq \leq t\mathbb{1} \\ &\quad\ \ \mathbb{1}^T q = 1 \\ &\quad\ \ q \geq \vec{0} \end{aligned}$$

Convert $\mathbb{1}^T p = 1$ to $\mathbb{1}^T p \leq 1$ and substitute $\vec{y} = \frac{1}{t}\vec{q}$:

$$\begin{aligned} &\min\ t \\ &\text{s.t. } My \leq \mathbb{1} \\ &\quad\ \ \mathbb{1}^T q \leq 1/t \\ &\quad\ \ q \geq 0 \end{aligned} \iff \begin{aligned} &\max\ \mathbb{1}^T q \quad (= 1/t) \\ &\text{s.t. } My \leq \mathbb{1} \\ &\quad\ \ y \geq 0 \end{aligned}$$

So we have proved strong duality when $b = c = \mathbb{1}$.
By scaling individual variables separately, can generalize to $b, c \geq \vec{0}$.