Optimization Methods in Finance **(EPFL, Fall 2010)**

Lecture 3: The Minimax theorem & Online portfolio selection

6.10.2010

Lecturer: Prof. Friedrich Eisenbrand *Scribe: Yoav Zemel*

Proof of the Minimax Theorem

Reminder : Our goal is to prove

Theorem 1 (Minimax)**.** *Let $A \in \mathbb{R}^{m\times n}$ be a payoff matrix then*

$$\min_Q \max_P A(Q,P) = \max_P \min_Q A(Q,P), \qquad A(Q,P) = Q^T AP$$

We have shown that "$\geq$" is easy, so it is sufficient to prove direction "$\leq$".
We wish to prove the minimax theorem using the randomized weighted majority algorithm that was presented in earlier lectures. In the algorithm we assumed that the loss vectors are in $[0,1]^N$, so we need a reduction to this restricted case.

Lemma 2. *It is sufficient to prove the minimax theorem for the case $A \in [0,1]^{m\times n}$.*

Proof. Let $\kappa = \min_{i,j} A_{ij}$ and $B = A - \kappa$ then B has nonnegative values. By a slight abuse of notation, we use κ also for the matrix that has all entries κ. The case $B = 0$ is a trivial, hence we may let $\gamma = \max_{i,j} B_{ij} > 0$ and $C = B\frac{1}{\gamma}$, then $C \in [0,1]^{m\times n}$. For P,Q with $\sum_i P_i = \sum_j Q_j = 1$ we have

$$A(Q,P) = (\gamma C + \kappa)(Q,P) = \gamma C(Q,P) + \kappa(Q,P) = \gamma C(Q,P) + \kappa$$

and the lemma follows. □

It is always more convenient to deal with finitely many objects than with infinitely many. In order to do so, we need to introduce the notion of best response. Given a choice of one player, what is (or are) the best choice (or choices) for the other player ?

Definition 3. A **BEST RESPONSE** to a mixed row strategy Q^* is a mixed column strategy $\tilde{P}$ that maximizes $(Q^*)^T AP$, where P ranges over mixed column strategies.
Similarly, a best response to a mixed column strategy P^* is a mixed row strategy $\tilde{Q}$ that minimizes $Q^T AP^*$, where Q ranges over all mixed row strategies.

It is clear that there is a best response, because we maximize or minimize a continuous function over a compact non-empty set. Moreover, there is always a simple one. This will give us finiteness.

Lemma 4. *There exists a pure best response (that is, a best response which is a pure strategy).*

Proof. We prove that there always exists a pure best response for a row strategy. The case of a column strategy is completely analogous.

Let Q^* be a fixed row strategy and $\tilde{P}$ be a best response of the column player. The $\tilde{P}_j$'s are nonnegative and sum up to one, so the expected payoff is

$$(Q^*)^T A\tilde{P} = \sum_{j=1}^{n}((Q^*)^T A)_j \tilde{P}_j \leq \sum_{j=1}^{n} \max_k((Q^*)^T A)_k \tilde{P}_j =$$

$$= \max_k((Q^*)^T A)_k = (Q^*)^T A e_{k'}, \qquad k' = \operatorname*{argmax}_k((Q^*)^T A)_k$$

where $e_{k'}$ is the vector with $e_{k'}(i) = 1$ if $i = k'$ and 0 otherwise. Hence $e_{k'}$ is the desired pure best response. □

We can therefore restate the minimax theorem :

$$\min_Q \max_j A(Q, j) = \max_P \min_i A(i, P)$$

where Q is a mixed row strategy, and j ranges only over the pure column strategies. Similarly, P is a mixed column strategy and i ranges only over the pure row strategies.

We are now ready to use the weighted majority algorithm in order to prove the minimax theorem. We need to define who (or what) are the experts, what is the loss vector, and what is the update rule. We use the algorithm in the following setting : The experts are the rows of A. At time t we have a probability distribution p^t on the experts, so p^t can actually be considered as a mixed row strategy. At time t nature chooses the loss vector ℓ^t to be the column of A that maximizes the expected loss of Ralf if he plays strategy p^t. In other words, the column ℓ^t is A^{j_t} where $j_t = \operatorname{argmax}_j (p^t)^T A^j$. That is, the loss vector is one of Celine's pure best responses for the row strategy p^t.

Suppose that there are N rows, that is, N experts. Let $\delta > 0$ be small and $T \geq \frac{4\ln N}{\delta^2}$, and set $\varepsilon = \frac{\delta}{2}$. The update rule is $w_j^{t+1} = w_j^t(1 - \varepsilon\ell_j^t)$. After T steps, the output row strategy that the algorithm returns is the average of the probability distributions $R = \frac{p^0 + \cdots + p^{T-1}}{T}$ (note that R is indeed a strategy, since $\sum_j R_j = \frac{T}{T} = 1$). We now show that R is almost the best strategy.

We have seen that the average expected loss of the forecaster in the randomized weighted majority algorithm after $O(\frac{\ln N}{\delta^2})$ rounds is about the loss of the best expert. More formally, we have proven that

Theorem 5. *Under the setting of the weighted majority with the update rule* $w_j^{t+1} = w_j^t(1 - \varepsilon\ell_j^t)$ *we have*

$$\frac{E[\widehat{L}]}{T} \leq \frac{L^j}{T} + \delta, \qquad T \geq \frac{2\ln N}{\varepsilon\delta}, \varepsilon \leq \min\left\{\frac{\delta}{2}, \frac{1}{2}\right\}.$$

So for $R = \frac{p^0 + \cdots + p^{T-1}}{T}$, and by Lemma 4, we have

$$\min_Q \max_P A(Q,P) = \min_Q \max_j A(Q, j) \leq \max_j A(R, j) = \frac{1}{T} \max_j \sum_{t=0}^{T-1} A(p^t, j)$$

since the minimum on all possible Q's cannot exceed what the row strategy R gives. Now the right hand side would have been bigger if we allow choosing j for each time separately, so

$$= \frac{1}{T}\max_j \sum_{t=0}^{T-1} A(p^t, j) \leq \frac{1}{T}\sum_{j=0}^{T-1} \max_j A(p^t, j) = \frac{1}{T}\sum_{j=0}^{T-1} A(p^t, j_t)$$

But $A(p^t, j_t)$ is exactly $E[\widehat{L}^t]$. So by the simple observation that $\ell^t_j = A(i, j_t)$ and by Theorem 5,

$$= \frac{1}{T}\sum_{j=0}^{T-1} E[\widehat{L}^t] = \frac{E[\widehat{L}]}{T} \leq \min_i \frac{\sum_{t=0}^{T-1} A(i, j_t)}{T} + \delta = \min_i A(i, P^*) + \delta \leq \max_P \min_i A(i, P) + \delta$$

where P is the column strategy $P^* = \frac{1}{T}\sum_{t=0}^{T-1} e_{j_t}$. We may let $\delta \to 0$, and the theorem is beautifully proved.

Example 6. *Recall our example from lecture two*

	A	*B*	*Money market*
Up	2	1.5	*1*
Stable	*1.2*	*1.7*	*1.3*
Down	*0.8*	*1.2*	*1.4*

After scaling (simply dividing by 2 in this case), the matrix becomes

$$A = \begin{pmatrix} 1 & 0.75 & 0.5 \\ 0.6 & 0.85 & 0.65 \\ 0.4 & 0.6 & 0.7 \end{pmatrix}$$

We set $\delta = 0.2$ *and* $\varepsilon = 0.1$.
We now try to find an optimal strategy using the algorithm.
At time 0, the weights are $(1,1,1)^T$ *so* $p^0 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})^T$. *So* $\ell^0 = (0.75, 0.85, 0.6)^T$ *for this is the best response of Celine to the strategy* p^0. *Use the update rule* $1 - \varepsilon e^j_t$, *so*

$$W^1 = \begin{pmatrix} 1 - 0.1 \cdot 0.75 \\ 1 - 0.1 \cdot 0.85 \\ 1 - 0.1 \cdot 0.6 \end{pmatrix}$$

Exercise : Repeat this 100 times and output the average distribution.

Online portfolio selection using the randomized weighted majority algorithm

We now go back to deal with finance. The setting is described in (Helmbold, Shapire, Singer & Warmouth : online portfolio selection using multiple updates, Mathematical Finance, vol 8, N4,

page 325-347).
Consider a portfolio with N stocks. The performance hinges on price relatives $(y_1^t, \dots, y_N^t)$ where y_i^t is

$$y_i^t = \frac{\text{price of stock } i \text{ at time } t+1}{\text{price of stock } i \text{ at time } t}$$

The portfolio is described by $x = (x_1, \dots, x_N)$ with $\sum_{i=1}^N x_i = 1$ and $x_i \geq 0$. Here x_i encodes the proportion of value in stock i. We use the algorithm to find an investment strategy and change x with time, such that we compare ourselves to the best constant investment strategy. It is clear that we cannot compete with the best strategy at all. In order to keep x constant, the player must balance the amounts of the stocks, regarding their prices.
The set of portfolios is the N-th dimensional simplex

$$\Sigma^N = \{(x_1, \dots, x_N) : \sum_{i=1}^N x_i = 1, x_i \geq 0\}$$

A portfolio where x remains the same over time is ***constant-rebalanced***. Note that such a trading strategy might involve lots of trading, which lead to transaction costs - the costs of selling/buying stocks. In our model we ignore the transaction costs.

Example 7. *Suppose that stock 1 always has return 1, while stock 2 has returns $\frac{1}{2}, 2, \frac{1}{2}, 2, \dots$. If we do not rebalance at all, we can only get a factor of two : Neither investment alone can increase by more than a factor of 2, even for infinite time.*
Consider the constant-rebalanced $x = (\frac{1}{2}, \frac{1}{2})$. What is the return ? On odd days, it is $\frac{1}{2}\frac{1}{2} + \frac{1}{2}1 = \frac{3}{4}$, and on even days it is $\frac{1}{2}2 + \frac{1}{2}1 = \frac{3}{2}$, so after two consecutive days the return is $\frac{9}{8} > 1$. After $2n$ days the return is $\left(\frac{9}{8}\right)^n$, so we have an exponential growth.

The motivating question we have is : Can we compete with the best constant-rebalanced portfolio in hindsight ? The answer is positive.

Let $y^0, \dots, y^{T-1}$ be the price relatives, and x^t the portfolio at time t. The return at time t is $(x^t)^T y^t$, or for short $x^t y^t$. Over the entire time horizon, it is

$$\prod_{t=0}^{T-1} x^t y^t \tag{0.1}$$

The best constant-rebalanced portfolio is some x that maximizes 0.1. This is of course equivalent to minimizing

$$\frac{1}{T} \sum_{t=0}^{T-1} - \ln(x^t y^t) \tag{0.2}$$

Our goal is to prove

Theorem 8. *One can find an online strategy $x_0, \dots, x_{T-1}$ such that*

$$\frac{1}{T} \sum_{t=0}^{T-1} \left(\ln(y^t x^*) - \ln(y^t x^t) \right) \leq 4\rho \sqrt{\frac{\ln N}{T}}$$

for any $x^* \in \Sigma^N$, where ρ is a bound on $\frac{y_i^t}{y_j^t}$ for all i, j, t.

We wish to prove theorem 8 using the randomized weighted majority algorithm. However, we have uncountably many ”experts” x. In order to deal with this problem, we use convexity.

Definition 9. $f : G \subseteq \mathbb{R}^n \to \mathbb{R}$ is **CONVEX** if $G = dom(f) \subseteq \mathbb{R}^n$ is convex, and if for any $x, y \in G, \lambda \in [0,1]$ we have

$$f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda) f(y)$$

Roughly speaking, $f : \mathbb{R} \to \mathbb{R}$ is convex if for any two points p_1 and p_2 on its graph, the line $[p_1, p_2]$ lies above the graph. This always happens if $f'' > 0$, for example.

Example 10. *Let $f^t : \sum^N \to \mathbb{R}$ by defined by*

$$f(x) = -\ln y^t x$$

then f^t is convex, because $-\ln x$ is a convex function.

The first order condition for convexity is

Lemma 11. *Suppose that $f : G \subseteq \mathbb{R}^n \to \mathbb{R}$ is differentiable and G is convex, then f is convex if and only if*

$$f(y) \geq f(x) + (\nabla(f(x)))^T (y - x)$$