EECS598: Prediction and Learning: It's Only a Game Fall 2013

Lecture 14: Game-theoretic Probability in Finance

Prof. Jacob Abernethy & Dr. Rafael Frongillo *Scribe: Murray (Yanzhe) Lei*

Announcements

- Homework due on Monday, Oct 28^{th}, 2013
- Guest lecturer Rafael Frongillo today.

1 Online Convex Optimization: A Gradient Descent Approach

1.1 General Framework of Online Convex Optimization (OCO)

Assume we have a *decision space* $X \subset \mathbb{R}^n$, which is convex, closed and compact. The game goes on such that:

For $t = 1, \ldots, T$

- Player plays some $x_t \in X$.
- Nature reveals some $l_t : X \to \mathbb{R}$ which is convex.
- Player suffers loss of $l_t(x_t)$.

Our goal is to minimize the regret regards the best static decision in hindsight:

$$Regret_T := \sum_{t=1}^{T} l_t(x_t) - \min_{x \in X} \sum_{t=1}^{T} l_t(x) \tag{1.1}$$

1.2 Online Gradient Descent (OGD) Approach

This algorithm is introduced by [Zin03]. Assuming $\{l_t(\cdot)\}_{t=1}^{T}$ are differentiable:

Starting with some arbitrary $x_1 \in X$.

For $t = 1, \ldots, T$

- $z_{t+1} \leftarrow x_t - \eta \nabla l_t(x_t)$.
- $x_{t+1} \leftarrow \Pi_X(z_{t+1})$, where $\Pi_X(\cdot)$ is the projection function.

The performance of *OGD* is described as follows.

Theorem 1.1. *If there exists some positive constant G, D such that*

$$\| \nabla l_t(x) \|_2 \leq G, \forall t, x \in X \tag{1.2}$$

$$\|X\|_2 ;= \max_{x,y \in X} \|x - y\|_2 \leq D \tag{1.3}$$

Then $Regret_T(OGD) \leq \mathcal{O}(DG\sqrt{T})$.

Proof. As usual we define potential function $\Phi_t := -\frac{1}{2\eta}\|x_t - x^*\|_2^2$. Notice that by (1.2), for every t we have:

$$\begin{aligned} \Phi_{t+1} - \Phi_t &\geq \nabla l_t(x_t)(x_t - x^*) - \frac{\eta}{2}\| \nabla l_t(x_t)\|_2^2 \\ &\geq l_t(x_t) - l_t(x^*) - \frac{\eta}{2}\| \nabla l_t(x_t)\|_2^2 \geq (l_t(x_t) - l_t(x^*)) - \frac{\eta}{2}G^2. \end{aligned}$$

Sum the above inequality up from $t = 1$ to T and by (1.1) we have

$$\Phi_{T+1} - \Phi_1 = \sum_{t=1}^{T} \Phi_{t+1} - \Phi_t \geq \sum_{t=1}^{T} (l_t(x_t) - l_t(x^*)) - \frac{\eta}{2}G^2 = Regret_T(OGD) - \frac{\eta G^2 T}{2}.$$

On the other hand,

$$\Phi_{T+1} - \Phi_1 = \frac{1}{2\eta}(\|x_1 - x^*\|_2^2 - \|x_{T+1} - x^*\|_2^2) \leq \frac{\|x_1 - x^*\|_2^2}{2\eta} \leq \frac{D^2}{2\eta}$$

To sum up we have

$$Regret_T(OGD) \leq \frac{\eta G^2 T}{2} + \frac{D^2}{2\eta} \sim \mathcal{O}(DG\sqrt{T})$$

and the last equation is achieved by choosing $\eta = D/(G\sqrt{T})$. QED.

Comment .

- What if the gradient doesn't exist? Use sub-gradient.
- The bound may not be optimal! Think about the expert setting, Theorem 1.1 only gives us $\mathcal{O}(\sqrt{nT})$.

2 Game-theoretic Probability in Finance

* This part of lecture is based on the first three chapters of [SV05], all the figures except Figure 4 are from [SV05].

2.1 Preliminary Examples

We consider a non-realistic stock that today sells for \$8 a share and tomorrow will either go down in price to \$5 or go up in price to \$10. There is a derivative x whose pay-off depends on tomorrow's stock price. The model is shown in Figure 1.

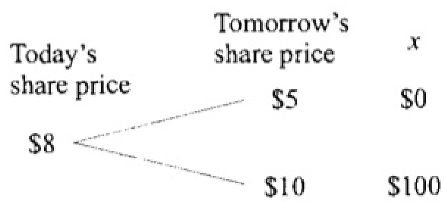

Figure 1: First Example

The question is: what would the value of x be if we have no arbitrage constraint? By "no arbitrage" we means no matter what happens there is no guaranteed positive pay-off. In this simple example, the no-arbitrage value of x would be \$60: we can buy 20 shares of the stock. If the stock goes down from \$8 to \$5, the loss of \$3 per share wipes out the \$60; if the stock goes up from \$8 to \$10, the gain of \$2 per share is just enough to cover the \$40 needed to provide x's pay-off \$100. If we set the value of \$x lower than \$60, then there is an arbitrage of buy more shares than 20, and similarly when we set the value higher then just buy less than 20 shares.

The no-arbitrage price could also be an range of price. Consider the second example in Figure 2.

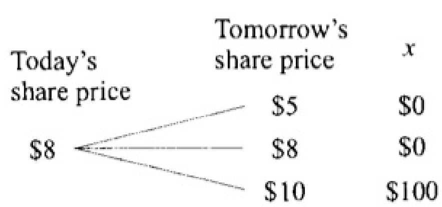

Figure 2: Second Example

Here we add another potential outcome of the stock price of tomorrow, which is remain at \$8. The pay-off of this outcome of derivative x' would be \$0. It is easy to see that in this case, the no-arbitrage price is a interval $[\$0, \$60]$.

The no-arbitrage price also exists in multi-stage game setting. Consider the third example in Figure 3

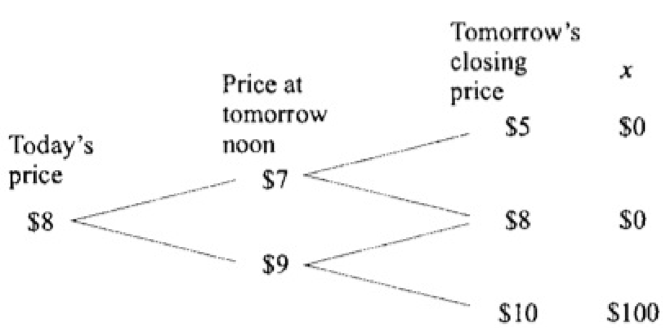

Figure 3: Third Example

Here we add one more stage to the second stage. The decision maker has a chance to adjust his number of share at the second stage after he observed the outcome of price at tomorrow noon ($7 or $9). The pay-off of the derivative x only depends on the final outcome, no matter what history the outcome has. In this example, the no-arbitrage price is $25. To hedge this price, we first buy 25 share of the stock today. We adjust this hedge at noon tomorrow, either by selling these shares (if the price goes down to $7) or by buying another 25 shares (if the price goes up to $9).

2.2 Game Theoretic Probability Theory

Now we formalize the discussed examples.

An outcome tree (like the ones above) is defined as the $World$, and the decision maker/gambler is defined as a $Skeptic$. The moves available to World may depend on moves he has previously made. But we assume that they do not depend on moves Skeptic has made. Skeptic's bets do not affect what is possible in the world, although World may consider them in deciding what to do next.

Each nodes in the tree represents a t. The initial situation is denoted by $\square$. We define the strategy function $\mathcal{P}(\cdot)$ as a real-valued function that maps a situation t to a decision. In the previous examples, t would be the stock prices and $\mathcal{P}(t)$ is the number of shares to buy.

We further define a *capital process* $\mathcal{K}^{\mathcal{P}}(t)$ as the Skeptic's capital in situation t. We assume $\mathcal{K}^{\mathcal{P}}(\square) = 0$, and Skeptics change of capital is linear, i.e. if the World change from situation t by ω, then Skeptics change of capital is $\mathcal{P}(t)\omega$, thus his capital at situation $t\omega$ is $\mathcal{K}^{\mathcal{P}}(t\omega) = \mathcal{K}^{\mathcal{P}}(t) + \mathcal{P}(t)\omega$.

We define a *martingale* as a function $s(t) = \alpha + \mathcal{K}^{\mathcal{P}}(t)$, where α is some *initial endowment*. We define a $variable$ $x(t)$ as a real-valued function that maps a terminal situation t to some monetary pay-off, which is the derivative in the previous examples. Given a variable x, we define its *upper expectation* and *lower expectation* as follows:

$$\overline{\mathbb{E}}[x] = \inf\{\alpha | \text{ exists a strategy } \mathcal{P} \text{ such that for all terminal situation } t, \alpha + \mathcal{K}^{\mathcal{P}}(t) \geq x(t)\}$$
$$\underline{\mathbb{E}}[x] = \sup\{\alpha | \text{ exists a strategy } \mathcal{P} \text{ such that for all terminal situation } t, \alpha + \mathcal{K}^{\mathcal{P}}(t) \leq x(t)\}$$

The upper expectation can be interpreted as the lowest amount of money at which the Skeptic can buy the derivative x, and the lower expectation can be interpreted as the highest amount of money at which the Skeptic can sell the derivative x. We can also write the expectations in a *regret* form:

$$\overline{\mathbb{E}}[x] = \inf_{\mathcal{P}} \sup_{t} \{x(t) - \mathcal{K}^{\mathcal{P}}(t)\}$$
$$\underline{\mathbb{E}}[x] = \sup_{\mathcal{P}} \inf_{t} \{x(t) - \mathcal{K}^{\mathcal{P}}(t)\}$$

If $\overline{\mathbb{E}}[x] = \underline{\mathbb{E}}[x]$, we define the *expectation* as $\mathbb{E}[x] = \overline{\mathbb{E}}[x] = \underline{\mathbb{E}}[x]$.

With the definition of expectation we can further define probability. Similarly we have the following definition for *upper probability* and *lower probability*:

$$\overline{\mathbb{P}}[E] = \overline{\mathbb{E}}[\mathbb{I}_E], \qquad \underline{\mathbb{P}}[E] = \underline{\mathbb{E}}[\mathbb{I}_E]$$

where the *event* E is a set of terminal situation and

$$\mathbb{I}_E(t) = \begin{cases} 1, & if\, t \in E \\ 0, & o.w. \end{cases}$$

Similarly, we define *probability* as $\mathbb{P}[E] = \overline{\mathbb{P}}[E] = \underline{\mathbb{P}}[E]$ if the latter two are equal.

One way to visualize this formulation is through a $2-D$ coordination system. Consider all the (t,x) pairs in previous example two. The strategy that could buy such a derivative x would be any line that covers all three points from above, and the upper expectation and upper probability would be the intercept and slope of the line. The similar hold for the lower expectation and lower probability. The expectation and probability exists if and only if the line that covers all the points from above coincide with the line that covers from below. In Figure 4 it is the dashed red line.

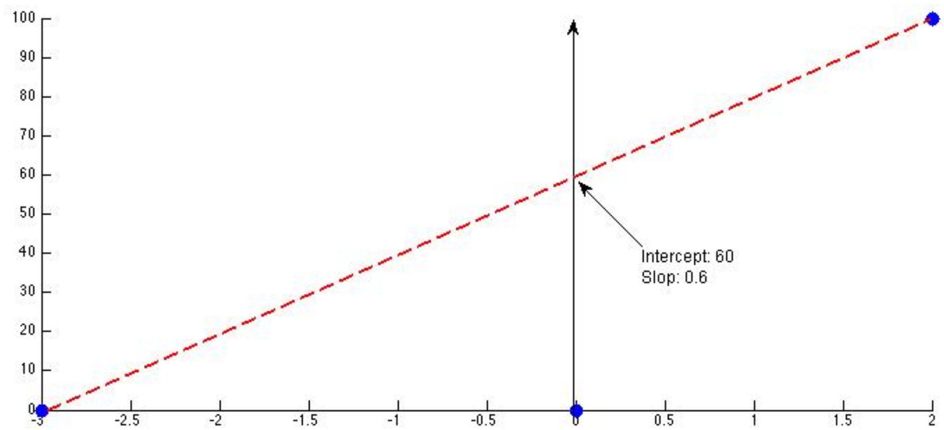

Figure 4: Interpretation of probability and Expectation

After we define expectation and conditional expectation we can revisit some obvious property of martingale:

$$\mathbb{E}_{\omega}[s(t\omega)|t] = s(t)$$

This is similar as the standard measure-base probability theory. It says your expected capital over all possible change ω given that your are in situation t is exactly your capital at situation t, i.e. $s(t)$.

2.3 Strong Law of Large Number

Consider a Skeptic with initial endowment $\alpha = 1$. The World plays in a infinite time scale $i = 1, \ldots, N, \ldots$. At round i, the World chooses $\omega_i \in [-1,1]$. Denote the *path* t as $t = \omega_1\omega_2\omega_3\ldots$. We stated SLLN without prove it.

Theorem 2.1. *If we define a event* $E := \{t | \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} \omega_i = 0\}$, *then* $\overline{P}(\neg E) = 0$

References

[SV05] Glenn Shafer and Vladimir Vovk. *Probability and finance: it's only a game!*, volume 491. John Wiley & Sons, 2005.

[Zin03] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. 2003.