CS281B/Stat241B (Spring 2008) Statistical Learning Theory **Lecture: 27**

Universal Portfolios

Lecturer: Peter Bartlett *Scribes: Boriska Toth and Oriol Vinyals*

1 Portfolio optimization setting

Suppose we have m instruments (i.e. stocks). At each period of time t, we have wealth V_t, and we invest a proportion $p_{t,i}$ in each instrument i. Instrument i increases in value by a factor $x_{t,i} \geq 0$. Our capital at period t thus increases by a factor $\frac{V_t}{V_{t-1}} = \sum_i \ x_{t,i}p_{t,i} = x_t'p_t$.

p_t is in the simplex Δ_m, and x_t is a random m-dimensional vector.

The first question to consider is: what is a fair bet? Is $\mathbb{E}\ x'b$ the right measure? Should one be satisfied making an investment of b, when $\mathbb{E}\ x'b > 1$?

In this setting of investment, one cares greatly about *risk*, in other words, the variance and not just the expectation of returns. The famous example below demonstrates.

Example. St. Petersburg paradox

Suppose x is a r.v. such that, at each time $t = 1, 2, ..$ there is a coin flip. Let $\tilde{t}$ denote the first time a "heads" comes up. Then the payout is $2^{\tilde{t}}$ for $t = \tilde{t}$, and 0 otherwise.

Then

$$\mathbb{E} \sum_{t=1}^{\infty} 2^i 2^{-i} = \sum_{t=1}^{\infty} 1 = \infty$$

Even though this procedure has infinite expected return, it doesn't make sense to invest much money in it, as there is very low probability of a large return. (Note that here we are using an additive and not multiplicative model of increases in wealth at each time.)

Four main approaches to portfolio optimization, that all capture the tradeoff between expected wealth and risk, are:

1. Markovitz approach
2. Utility function approach
3. Optimizing for optimal growth rate
4. Constant rebalanced portfolio (CRP) approach

1.1 Markovitz optimization

This approach was developed by Markovitz in the 50's. We have returns $X \in \mathbb{R}^m_+$, with $\mathbb{E}X = \mu$ and $\mathrm{Var}(X) = \Sigma$.

We maximize expected return, subject to a constraint on the risk (this example doesn't contain multiple time steps):

$$\max_{p \in \Delta_m} \ \mathbb{E}X : \quad \mathrm{Var}(X'p) \leq R \qquad \Leftrightarrow \qquad \max_{p \in \Delta_m} \ \mu' p : \quad p'\Sigma p \leq R$$

Without the constraints that bound variance, we have a linear optimization [max $\mu' p$] over a simplex Δ_m. The optimal solution to that is always to pick a vertex of the simplex, which corresponds to the pure strategy of investing in one instrument. The constraint $p'\Sigma p \leq R$ restricts $p \in \Delta_m$ to also be within an ellipsoid, so that the solution might involve mixing over more than one instrument (diversification helps control risk).

1.2 Utility function approach

This approach consists of defining a utility function $U()$ of values, chosen to take into account the tradeoff between expected wealth and risk, so that we maximize $\mathbb{E}\ U(x'p)$. This is the solution to the St. Petersburg problem proposed by Bernoulli.

To capture the criterion that a given expected gain with higher variance is worth less than one with lower variance, we should use a concave function with

$$\mathbb{E}U(V) \leq U(\mathbb{E}V)$$

Taking the second order Taylor approximation to $U(V)$ at $V = \mathbb{E}V$, and taking the expectation of both sides, we have

$$\mathbb{E}U(V) \approx U(\mathbb{E}V) + \frac{1}{2}U''(\mathbb{E}V)\mathrm{Var}(V)$$

For a concave function, $U''(\mathbb{E}V) \leq 0$, so we have a term that penalizes variance in optimizing the expected utility.

The log() function is used widely as a utility function, as it is in the constant rebalanced portfolio approach.

1.3 Example using log utility function

Suppose we have 2 instruments, with one of them risk-free, such that at all times:

$$\Pr(X_{t,1} = 1) = 1$$

$$\Pr(X_{t,2} = 0) = p < \frac{1}{2}$$

$$\Pr(X_{t,2} = 2) = 1 - p$$

We start with $V_0 > 0$ capital. Putting all our capital on 2, we have $\mathbb{E}\ V_T = (2(1-p))^T V_0$. This goes to infinity in the limit of T.

Now suppose we want to maximize the asymptotic growth rate of the log-of-return. Denote by b_t the proportion of capital bet on instrument 2 at time t. Also define $W_t = 1[X_{t,2} = 2]$ ("win").

Then

$$V_{t+1} = V_t(1 + b_t)^{W_t}(1 - b_t)^{1-W_t}$$

The asymptotic growth rate of the log-of-return is:

$$G = \lim_{T\to\infty} \frac{1}{T} \log_2 \frac{V_T}{V_0} =$$

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} W_t \log_2(1+b_t) + (1-W_t)\log_2(1-b_t)$$

This equality comes from expanding $\frac{V_T}{V_0}$ as a telescoping product.

Note that if $Z_t \in Z,\ t = 1, 2, ..$ are i.i.d. random variables, and $f()$ a real function on Z, then with probability 1 we have

$$\lim_{T\to\infty} \tfrac{1}{T} \ \textstyle\sum_{t=1}^{T} f(Z_t) \quad = \quad \mathbb{E}(Z_t)$$

So we have that with probability 1

$$G \quad = \quad (1-p)\log_2(1+b_t) + p\log_2(1-b_t) \quad = \quad \mathbb{E}\log\frac{V_{t+1}}{V_t}$$

It is straightforward to see that the optimal investment at any time t is $b_t^* = 1 - 2p$.

Note that in this example, since we have i.i.d. returns, it doesn't actually make a difference in expectation to average growth rate over time. Both the expected growth rate, and the optimal b_t, are the same in each period.

Also note that there is a close connection between maximizing the expected log return, and finding the ideal channel capacity in information theory.

2 Constant rebalanced portfolios

Definition. For random returns $x_t \in \mathbb{R}_+^m$, a strategy b^* is **log optimal** if
$b_t^*(x_1,\ x_2,\ ...\ x_{t-1}) = \mathrm{argmax}_{b\in\Delta_m}\ \mathbb{E}[\log(b' x_t)|x_0,\ ...\ x_{t-1}]$.

To motivate this definition, suppose we want to maximize $\mathbb{E}$ $(\log \frac{V_T}{V_0})$ by choosing strategy $b_t^*(x_1,\ x_2,\ ...\ x_{t-1})$ that can depend on all known return values. Expressing $\frac{V_T}{V_0}$ as a telescoping product, and using $\frac{V_t}{V_{t-1}} = b_t' x_t$, we have

$$\max_{\{b_t\}} \ \mathbb{E}\ [\sum_{t=1}^{T} \log(b_t' x_t)]$$

Thus we choose each b_t as

$$b_t \quad = \quad \mathrm{argmax}_{b_t(x_1,\ ...\ x_{t-1})}\ \mathbb{E}\ [\log(b_t' x_t)] \quad = \quad \mathrm{argmax}_{b_t}\ \mathbb{E}\ [\log(b_t' x_t)|x_1,\ ...\ x_{t-1}]$$

Note that if the returns are i.i.d., then $b_t^*(x_1,\ x_2,\ ...\ x_{t-1}) = \mathrm{argmax}_{b\in\Delta_m}\ \mathbb{E}[\log(b' x]$, so we have a constant log optimal strategy b^* at each time. This is called a **constant rebalanced portfolio** (CPR). This means we must rebalance our investment, after the stocks have grown at non-uniform rates to yield a different balance than b^*, such that we use strategy b^* at each discrete time interval.

Theorem 2.1. Given i.i.d. returns x_t, if the log optimal strategy b^* has capital growth $V_0,\ V_1^*,\ ...\ V_T^*$ and some strategy b has growth $V_0,\ V_1,\ ...\ V_T$, then

$$\limsup_{T\to\infty}\ \frac{1}{T}\log\frac{V_T}{V_T^*}\ \le\ 0 \qquad \text{almost surely}$$

We can see that this makes sense because $\log\frac{V_T}{V_T^*} = \log\frac{V_T}{V_0} - \log\frac{V_T^*}{V_0} = \sum_{t=1}^{T}\left[\log(x_t'b) - \log(x_t'b^*)\right]$. We saw that $\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\ \log(x_t'b^*) \to \mathbb{E}\log(x'b)$, which is maximized by b^*.

Thus, in the i.i.d. case, choosing the log optimal portfolio b^* as a CPR is optimal.

To summarize, in the i.i.d. case, there is a constant strategy vector b, such that if the investment is rebalanced at each step to strategy b, then we have optimal asymptotic growth rate of log returns. This overall investment strategy of finding a log optimal investment vector, and rebalancing to keep investment constant, is called the CRP.

3 CRP_b

We define the CRP_b as the strategy described above of rebalancing a portfolio at each time step to the constant investment vector $b \in \Delta_m$ (here we will not just be considering the case of i.i.d. returns and log optimal b).

E.g. 1: Let

$$X_t = \left\{\begin{pmatrix}2\\ \frac{1}{2}\end{pmatrix}, \begin{pmatrix}\frac{1}{2}\\ 2\end{pmatrix}, \begin{pmatrix}2\\ \frac{1}{2}\end{pmatrix}, \dots\right\}$$

If we pick the constant rebalanced strategy of $b = \begin{pmatrix}\frac{1}{2}\\ \frac{1}{2}\end{pmatrix}$, that is, uniform allocation of money, the growth is exponential with rate=$\frac{5}{4}$. On the other hand, it is easy to check that any non-rebalanced strategy, where money is not moved between the two instruments after the first investment, achieves a growth of only 1.

E.g. 2: Dow Jones Industrial Averages (DJIA), which is a commonly used index to evaluate the US market performance and which equally distributes the weight among 30 companies, is another example of CRP. The year-to-date change for 2008 is $-1.56\,\%$.

Theorem 3.1. There is a strategy b_U (universal strategy) for which, for all $X_1, \dots, X_T$, $\log(V_T(b_U)) \ge \log(V_T(b^*)) - (m-1)\log(T+1) - 1$, where CRP_{b^*} is the CRP having optimal log returns (best constant strategy in retrospect).

Before giving proof, we first demonstrate a simplified version of this theorem. Consider the m pure CRPs: $CRP_b,\ b \in \{e_i,\ i = 1\dots m\}$. Suppose we want to do reasonably well (in the usual, log optimal sense) as compared to the best of these CRP's. A good strategy then is to use CRP_{b_W}, where b_W places $\frac{1}{m}$ in each instrument and does not rebalance. Then

$$\log(V_T(b_W)) = \log\left(\sum_{j=1}^{m}\prod_{t=1}^{T} X_{t,j}V_0/m\right) \ge \log\max_j\prod_{t=1}^{T} X_{t,j}V_0/m =$$

$$= \max_j \log\left(\prod_{t=1}^{T} X_{t,j}V_0\right) - \log m = \max_j \log(V_T(e_j)) - \log m$$

which means that, by not knowing which is the best instrument, we pay at most a price of $\log m$.

PROOF. The universal strategy b_U consists of allocating capital uniformly across the simplex Δ_m at the first time step, so that we have $V_0 d\mu(b)$ capital at each strategy $b \in \Delta_m$. Once we allocate capital uniformly across the simplex, money is never moved between the different portfolios b, so the different portfolios grow at different rates and don't have uniform allocation across the simplex after the first iteration. Note that the money invested in each portfolio $b \in \Delta_m$ is rebalanced at each step.

Further, let $V_t(b)$ denote the ratio of capital at time t to capital at time 0 resulting from an investment of CRP_b. By definition, $V_0(b) = 1 \ \forall b$. Thus, at time t, we have $V_t(b)\ V_0 d\mu(b)$ capital on strategy b.

The actual strategy chosen at each time t is found by integrating across the simplex as follows:

$$b_t = \frac{\int_{\Delta_m} bV_{t-1}(b)\ V_0 d\mu(b)}{\int_{\Delta_m} V_{t-1}(b)\ V_0 d\mu(b)} = \frac{\int_{\Delta_m} bV_{t-1}(b) d\mu(b)}{\int_{\Delta_m} V_{t-1}(b) d\mu(b)}$$

Let us now consider the set S a ball (in a topological sense) on the simplex centered on the optimal solution b^* and with radius ε,

$$S = \{(1-\varepsilon)b^* + \varepsilon a \ : \ a \in \Delta_m\}$$

For $b \in S$,

$$\frac{V_1(b)}{V_0} = \frac{V_1((1-\varepsilon)b^* + \varepsilon a)}{V_0} \geq (1-\varepsilon)\frac{V_1(b^*)}{V_0}$$

and, using the same inequality for times t, $t-1$ and telescoping the product, we get $\frac{V_T(b)}{V_T(b^*)} \geq (1-\varepsilon)^T$. So far, we've seen that choosing $b \in S$ is not a bad choice. Also, the proportion of V_0 allocated to S is $\mu(S) = \mu(\{\varepsilon a \ : \ a \in \Delta_m\}) = \varepsilon^{m-1}$. It is straightforward to see that

$$\log \frac{V_T(b_U)}{V_T(b^*)} \geq \log((1-\varepsilon)^T \varepsilon^{m-1})$$

If we choose $\varepsilon = \frac{1}{T+1}$, $\log(V_T) \geq \log(V_T(b^*)) - (m-1)\log(T+1) - 1$. $\square$

Some references here are Kelly, Breiman, Algoet, and Cover.

4 Prediction with log loss

Given a distribution X_t, we want to predict it using an estimation $\widehat{p_t}$. The log loss is defined as:

$$l(\widehat{p_t}, X_t) = -\log(\widehat{p_t}(X_t))$$

In the i.i.d. case, the expected log loss and the KL divergence are closely related (they are the same except for an additional term of the entropy of the true distribution).

If we have several predictions of the probability function $\widehat{p_{t,1}}, \ldots, \widehat{p_{t,m}}$, we can formulate the problem to find the optimal mixture weight in terms of log loss (eg. language models). In particular, we want to find the optimal $b \in \Delta_m$ to minimize the following loss function:

$$l(\widehat{p_t}, X_t) = -\log\left(\sum_{i=1}^{m} b_i \widehat{p_{t,i}}(X_t)\right)$$

This is equivalent to the problem presented in this lecture and, therefore, we can use portfolio strategies to solve it.