Gilles Stoltz

Prediction with Expert Advice: Theory and Applications.

Lecture #4.

→ Setting:

- N assets (Alitalia, Coca-cola, Google, etc.), $j=1,\dots N$
- We describe their evolution at day t by multiplicative factors x_{jt}, $j=1,\dots N$ $\rightarrow$ $\underline{x}_t = (x_{1t}, \dots x_{Nt}) \in \mathbb{R}_+^N$ is the market vector at day t; $x_{jt}=0$ possible...
- Investment at a certain day is given by a portfolio $\underline{P}$: $\underline{P} = (P_1, \dots P_N)$ is an element of the N-th order simplex $\mathcal{X}$ (that is, a probability distribution over $\{1,\dots N\}$)
- Portfolio $\underline{P}$ gets richer by a multiplicative factor of
$$\underline{P}\cdot\underline{x}_t = \sum_{j=1}^{N} P_j x_{jt}$$
at day t. This quantity is called the wealth factor of P at day t.

Def: [Investment strategy:] An investment strategy is a sequence $(\underline{Q}_t)_{t\geq 1}$ functions $\underline{Q}_t : (\mathbb{R}_+^N)^{t-1} \longrightarrow \mathcal{X}$

Not: • We denote $\underline{x}_1^{t-1} = (\underline{x}_1, \dots \underline{x}_{t-1})$ the first $t-1$ market vectors, and $\underline{Q}_t = \underline{Q}_t(\underline{x}_1^{t-1}) = \left(Q_{1t}(\underline{x}_1^{t-1}), \dots, Q_{Nt}(\underline{x}_1^{t-1})\right)$.

• The wealth factor achieved by $(\underline{Q}_t)_{t\geq 1}$ after n days of trading equals
$$\hat{S}_n = S_n\left((\underline{Q}_t), \underline{x}_1^n\right) = \prod_{t=1}^{n} \left(\sum_{j=1}^{N} Q_{jt}(\underline{x}_1^{t-1})\, x_{jt}\right).$$

Remark: Two constraints of real stock market investment are ignored:

At the end of day $(t-1)$, our wealth is distributed according to the distribution

$$\underline{Q}'_{t-1} = \left(\frac{Q_{1,t-1}\, x_{1,t-1}}{\underline{Q}_{t-1} \cdot \underline{x}_{t-1}}, \ldots, \frac{Q_{N,t-1}\, x_{N,t-1}}{\underline{Q}_{t-1} \cdot \underline{x}_{t-1}} \right)$$

and we want to rebalance to $\underline{Q}_t$ at the begining of day t:

(1) we assume we can always sell or buy as much as we want for each asset;

(2) and that this can be done without transaction costs.

Assumption (1) is reasonable if we are small traders; assumption (2) can be removed by taking into account some penalty factors for rebalancing (see Blum & Kalai '97).

<u>Example</u>: The simplest strategy is buy-and-hold: do not trade, just let it go: $\hat{S}_n = \frac{1}{N} \sum_{j=1}^{N} \prod_{t=1}^{n} x_{jt}$,

obtained with
$$\underline{Q}_t(\underline{x}_1^{t-1}) = \frac{\frac{1}{N}\sum_{i=1}^{N} \left(\prod_{s=1}^{t-1} x_{is}\right) \delta_i}{\frac{1}{N}\sum_{j=1}^{N} \prod_{s=1}^{t-1} x_{js}} \qquad \text{where } \delta_i = (0, \ldots 0, \underset{\substack{\uparrow \\ i\text{-th position}}}{1}, 0, \ldots 0)$$

Note that $0 \leqslant \log \dfrac{\max_{j=1,\ldots N} \prod_{t=1}^{n} x_{jt}}{\hat{S}_m} \leqslant \ln N.$

(Regret bounds will be of this form.)

$\longrightarrow$ <u>Assessment</u> of a strategy: the comparison class

The comparison class is formed by:

<u>Def</u>: [constantly rebalanced portfolios - CRPs:] The CRP $\underline{P}$, indexed by $\underline{P} \in \mathcal{X}$, uses the investment strategy $(\underline{Q}_t)_{t\geqslant 1} = (\underline{P})_{t \geqslant 1}$, that is, rebalances every day to $\underline{P}$: $\underline{Q}_t(\underline{x}_1^{t-1}) = \underline{P} \quad \forall \underline{x}_1^{t-1}$.

<u>Ex</u>:
- Putting all the money on one stock (say j) and letting it sit is a CRP investment strategy, with $\underline{P} = \delta_j$
- Buy-and-hold cannot be described by a CRP.

The wealth factor of the CRP $\underline{P}$ is

$$S_n(\underline{P}) = \prod_{t=1}^{n} \left(\sum_{j=1}^{N} P_j x_{jt} \right).$$

It can be big even on average markets.

Ex:
$$\underline{x}_1 = \underline{x}_3 = \underline{x}_5 = \dots = \begin{pmatrix} 1 \\ 1/2 \end{pmatrix}$$
$$\underline{x}_2 = \underline{x}_4 = \underline{x}_6 = \dots = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

Single stocks (n even): $\quad S_n(\delta_1) = S_n(\delta_2) = 1$

CRP $\begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}$: $\quad S_n\left(\begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}\right) = \left(\left(\tfrac12 + \tfrac12 \times \tfrac12\right)\left(\tfrac12 + \tfrac12 \times 2\right) \right)^{n/2}$

$$= \left(9/8\right)^{n/2} \longrightarrow \text{exponential growth.}$$

That CRPs are interesting strategies is summarized in the results below. (They use stochastic assumptions only to justify the interest of CRPs; we will not use stochasticity elsewhere.)

Lm: If markets vectors are given by an iid process $\underline{X}_1, \underline{X}_2, \dots$ then

$$\max_{\underline{P}} E\left[\log S_n(\underline{P})\right] = \max_{(Q_t)_{t \geq 1}} E\left[\log \hat{S}_n\right].$$

Proof:
$$E \log \hat{S}_n = \sum_{t=1}^{n} \underbrace{E\left[\log Q_t(\underline{X}_1^{t-1}) \cdot X_t\right]}$$

$$\begin{aligned}
&E\left[\log Q_t(\underline{X}_1^{t-1}) \cdot \underline{X}_t\right] \\
&= E\left[E\left[\log Q_t(\underline{X}_1^{t-1}) \cdot \underline{X}_t \mid \underline{X}_1^{t-1}\right]\right] \\
&\leq E\left[\max_{\underline{P}} E\left[\log \underline{P} \cdot \underline{X}_t\right]\right] \\
&= \max_{\underline{P}} E\left[\log \underline{P} \cdot \underline{X}_t\right] = \max_{\underline{P}} E\left[\log \underline{P} \cdot \underline{X}_1\right]
\end{aligned}$$

and
$$E \log S_n(\underline{P}) = \sum_{t=1}^{n} E\left[\log(\underline{P} \cdot X_t)\right] = n\, E\left[\log(\underline{P}\, \underline{X}_1)\right].$$

Even stronger is the following theorem (proof omitted, see Cover & Thomas '91, chap. 15).

Th: For an iid market $\underline{X}_1, \underline{X}_2, \dots$ (or even an only stationary market), let $P^* \in \arg\max E[\log \underline{P} \cdot \underline{X}_1]$, then:

For all investment strategies $(\underline{Q}_t)_{t \geq 1}$,

with probability 1: $\overline{\lim} \frac{1}{n} \log \frac{\hat{S}_n}{S_n(P^*)} \leq 0$

We now construct strategies such that in particular,

$\forall \underline{x}_1, \dots \underline{x}_n, \dots$ $\overline{\lim} \frac{1}{n} \log \frac{\max_{\underline{P} \in \mathcal{X}} S_n(\underline{P})}{\hat{S}_n} \leq 0.$

the regret of the strategy $(Q_t)_{t \geq 1}$ ...

Formal definition of the REGRET.

Wealth ratios $\underline{x}_1, \underline{x}_2, \dots$ are now only described, not modelled any longer in a stochastic way.

We consider a worst-case assessment, with respect to CRB, using log weath ratios:

$$R_n = R_n((\underline{Q}_t)_t) = \sup_{\substack{\underline{x}_1 \dots \underline{x}_n \\ \in (\mathbb{R}_+^N)^n}} \log \frac{\max_{\underline{P} \in \mathcal{X}} S_n(\underline{P})}{\hat{S}_n}$$

Remark: [Connection with prediction with expert advice:]

$y_t = \underline{x}_t$, experts are single stocks: $\delta_1, \dots \delta_N$,

$\underline{Q}_t = \underline{Q}_t(\underline{x}_1^{t-1})$ is indeed a convex combination of the experts' advices $\delta_1 \dots \delta_N$

and the loss function is $\ell(\underline{p}, \underline{x}) = -\log \underline{p} \cdot \underline{x} = -\log \sum_{j=1}^N p_j x_j$.

Then,

$$R_n = \sup_{\underline{x}_1^n} \sum_{t=1}^n \ell(\underline{Q}_t, \underline{x}_t) - \min_{\underline{P} \in \mathcal{X}} \sum_{t=1}^n \ell(\underline{P}, \underline{x}_t).$$

$\longrightarrow$ A simple strategy linked with the <u>exponentially weighted average forecaster</u>

Lm: $\quad a, b > 0: \quad -\log a - (-\log b) \leq \dfrac{b-a}{a}.$

Proof: This is a matter of convexity of $-\log$ (the slopes inequality):

$$-\log a - (-\log b) \leq (-1/a)(b-a):$$

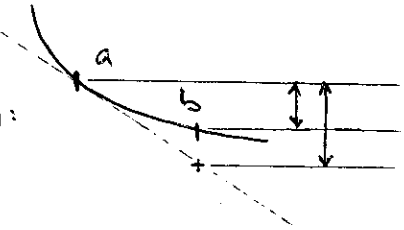

Can be seen also by using $\log(1+u) \leq u$ for all $u > -1$:

$$-\log a + \log b = \log \frac{b}{a} = \log\left(1 + \frac{b-a}{a}\right) \leq \frac{b-a}{a}.$$

Application of the lemma: upper bound on the regret linear in $\underline{P}$:

Fix $\underline{P}$ and $\underline{x}_1^n$:

$$\sum_{t=1}^{n} -\log(\underline{Q}_t \cdot \underline{x}_t) - \sum_{t=1}^{n} -\log(\underline{P} \cdot \underline{x}_t)$$

$$= \sum_{t=1}^{n} \Big((-\log \underline{Q}_t \cdot \underline{x}_t) - (\log \underline{P} \cdot \underline{x}_t) \Big)$$

$$\leq \sum_{t=1}^{n} \frac{1}{\underline{Q}_t \cdot \underline{x}_t} \left(\underline{P} \cdot \underline{x}_t - \underline{Q}_t \cdot \underline{x}_t \right)$$

$\uparrow$ linear in $\underline{P}$: $\quad \max\limits_{\underline{P}}$ "..." $= \max\limits_{\delta_1, \dots, \delta_N}$ "..."

Therefore, for all $\underline{x}_1^n$:

$$\log \frac{\max_{\underline{P} \in \mathcal{X}} S_n(\underline{P})}{\hat{S}_n} \leq \max_{j=1,\dots,N} \sum_{t=1}^{n} \frac{1}{\underline{Q}_t \cdot \underline{x}_t} \left(x_{jt} - \underline{Q}_t \cdot \underline{x}_t \right)$$

$$= \max_{j=1,\dots,N} \sum_{t=1}^{n} \sum_{k=1}^{N} Q_{kt} \, \tilde{\ell}_{kt} - \sum_{t=1}^{n} \tilde{\ell}_{jt}$$

where $\quad \tilde{\ell}_{kt} = -\dfrac{x_{kt}}{\underline{Q}_t \cdot \underline{x}_t}$

Say $0 < m \leq x_{jt} \leq M$ for all j and t (bounded market). Then, one can also take

$$\tilde{\ell}_{kt} = \frac{M}{m} - \frac{x_{kt}}{\underline{Q}_t \cdot \underline{x}_t} \in [0, M/m].$$

If $$Q_{kt} = \exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_{ks}\right) \Big/ \sum_{i=1}^{N} \exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_{is}\right)$$

then $$\sum_{t=1}^{n} \sum_{k=1}^{n} Q_{kt}\, \tilde{\ell}_{kt} - \sum_{t=1}^{n} \tilde{\ell}_{jt} \leq \frac{\ln N}{\eta} + \eta \left(\frac{M}{m}\right)^2 \frac{n}{8}$$

(cf. previous classes)

$$= \square \frac{M}{m} \sqrt{n \ln N}$$ for a proper choice of η

↳ the investment strategy thus defined is called EG.

Th: for EG, $$\sup_{\substack{\underline{x}_1, \dots \underline{x}_n \\ \in [m, M]^N}} \log \frac{\max_{\underline{P} \in \mathcal{X}} S_n(\underline{P})}{\hat{S}_n} \leq \square \frac{M}{m} \sqrt{n \ln N}$$

Rk:

- A modification of EG, using among others a mixing of the Q_t defined above with the uniform portfolio $(1/N, \dots 1/N)$ ensures that
$$R_n = \sup_{\underline{x}_1, \dots \underline{x}_n \in \mathbb{R}_+^N} \log \frac{\max_{\underline{P} \in \mathcal{X}} S_n(\underline{P})}{\hat{S}_n} \leq 10\, N\, n^{2/3}.$$

- The $\sqrt{n}$ rate for the regret of EG is unavoidable; however, as we shall see below, the faster $\ln n$ rate can be guaranteed. In the analysis we lost a lot when linearly upper bounding the (in $\underline{P}$) concave quantity R_n by its (in $\underline{P}$) linear upper bound. The max. is achieved in a corner of $\mathcal{X}$ in the upper bound, but in the interior of $\mathcal{X}$ in the regret R_n.

⟶ Cover's ('91) UNIVERSAL PORTFOLIO.

Toy example: It is easy to be competitive wrt a finite number of CRPs $P_1, \dots P_m$: just use buy-and-hold on these portfolios: divide the money in m parts on day 1, invest each fraction according to a P_j and let it at; we obtain $\hat{S}_n = \frac{1}{m} \sum_{j=1}^{m} S_n(\underline{P}_j)$

and $\log \frac{\max_{j=1..N} S_n(\underline{P}_j)}{\hat{S}_n} \leq \log N.$

This uses (see page 2):
$$\underline{Q}_t(x_1^{t-1}) = \frac{\frac{1}{N}\sum_{i=1}^{N} S_{t-1}(\underline{P}_i) \times \underline{P}_i}{1/N \sum_{j=1}^{N} S_{t-1}(\underline{P}_j)}.$$

<u>Goal</u>: Now, be competitive wrt all CRPs!
Take μ the uniform distribution on $\mathcal{X}$ (more on that later)
and use
$$\underline{Q}_t(x_1^{t-1}) = \frac{\int_{\mathcal{X}} S_{t-1}(\underline{P})\, \underline{P}\, d\mu(\underline{P})}{\int_{\mathcal{X}} S_{t-1}(\underline{P})\, d\mu(\underline{P})}$$

The wealth obtained by this strategy is
$$\hat{S}_n = \prod_{t=1}^{n} \underline{Q}_t \cdot \underline{x}_t = \prod_{t=1}^{n} \frac{\int_{\mathcal{X}} \underbrace{S_{t-1}(\underline{P}) \times \underline{P} \cdot \underline{x}_t}_{= S_t(\underline{P})}\, d\mu(\underline{P})}{\int_{\mathcal{X}} S_{t-1}(\underline{P})\, d\mu(\underline{P})}$$
$$= \int_{\mathcal{X}} S_n(\underline{P})\, d\mu(\underline{P}) \quad \text{by telescoping.}$$

<u>Th</u>: $R_n \leq (N-1)\log n + O(1).$

<u>Proof</u>: Let $\underline{P}^*$ be the optimal portfolio on a fixed sequence $\underline{x}_1, \dots \underline{x}_n$,
ie, $S_n(\underline{P}^*) = \max_{\underline{P} \in \mathcal{X}} S_n(\underline{P}).$

Portfolios of the form $\underline{P} = (1-\alpha)\underline{P}^* + \alpha z$ (for $z \in \mathcal{X}$)
achieve a wealth $S_n(\underline{P}) \geq (1-\alpha)^n S_n(\underline{P}^*).$

How many such portfolios do we have?
$$\mu\{\underline{P} : \exists z \in \mathcal{X} \text{ with } \underline{P} = (1-\alpha)\underline{P}^* + \alpha z\}$$
$$= \alpha^{N-1} \mu\mathcal{X} = \alpha^{N-1}$$

Thus, the wealth $\hat{S}_n$ is at least

$$\hat{S}_m = \int_{\mathcal{X}} S_n(\underline{P})\, d\mu(\underline{P}) \geqslant \alpha^{N-1} (1-\alpha)^m \, S_n(\underline{P}^*)$$
$$\geqslant \square \left(\frac{1}{m}\right)^{N-1} S_n(\underline{P}^*)$$

($\square$ is a positive lower bound on the sequence $(1-1/n)^n \xrightarrow[n\to+\infty]{} 1/e$)

by choosing $\alpha = 1/m$ for the analysis

Finally, $\forall\, \underline{x}_1, \dots \underline{x}_n$, $\log \dfrac{S_n(\underline{P}^*)}{\hat{S}_n} \leqslant (N-1)\log n + \log \square$.

<u>Remark:</u> • [For the practical implementation :] it is impossible in general to compute explicitely the $\int_{\mathcal{X}} \dots d\mu(\underline{P})$; even a finite grid does usually not do the job (# of points exponential in N).

One resorts rather to Monte-Carlo approximations :

$$\int_{\mathcal{X}} f(\underline{P})\, d\mu(\underline{P}) \approx \frac{1}{N} \sum_{j=1}^{N} f(\underline{P}_j)$$

where $\underline{P}_1, \dots \underline{P}_N$ drawn iid at random according to μ.

(A way to draw according to μ is to draw $X_1, \dots X_{N-1}$ iid at random in $[0,1]$, sort them : $X_{(1)}, \dots X_{(N-1)}$; and then $(X_{(1)}, X_{(2)} - X_{(1)}, \dots, X_{(N-1)} - X_{(N-2)}, 1 - X_{(N-1)})$ is distributed according to μ.)

• Ordentlich & Cover '98 proved that $N \ln n$ is the optimal order of magnitude of R_n.

• In practice : Results not extraordinary (even with the survivor bias) ... The good thing would be to use these methods not directly on stocks but on a finite number N' of base investment strategies (back to our experts!). All these techniques are indeed rather <u>meta-learning</u> techniques !