Mathematics of Finance, Courant Institute, Spring 2019
https://www.math.nyu.edu/faculty/goodman/teaching/MathFin2019/MathFinance.html
Always check the class message board before doing any work on the assignment.

Assignment 7, due April 23

Corrections: (computational exercise added)

1. **True/False**.
   (a) If X is uncorrelated with Y, then X^3 is uncorrelated with Y^3.
   (b) Suppose w and v are (portfolio weights for) efficient portfolios with the same assets that satisfy the same budget constraint $\mathbf{1}^t w = \mathbf{1}^t v = 1$. If $v \neq w$, and u is any other efficient portfolio with the same assets, then u is a linear combination of v and w.
   (c) If v and w are as in part (b) above, then any linear combination of v and w is an efficient portfolio.

2. **Multiple choice** Which of the following matrices is the covariance matrix of a pair of random variables (X, Y)?
   (a)
   $$C = \begin{pmatrix} 3 & 2 \\ 1 & 4 \end{pmatrix}$$
   (b)
   $$C = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}$$
   (c)
   $$C = \begin{pmatrix} -1 & 2 \\ 2 & 3 \end{pmatrix}$$

3. Suppose
   $$X = \begin{pmatrix} X_1 \\ X_2 \end{pmatrix}$$
   is a bi-variate random variable. The variables X_1 and X_2 have 2×2 covariance matrix C and means $\text{E}[X_1] = \mu_1$, $\text{E}[X_2] = \mu_2$. We call a non-random generic point
   $$x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$
   The mean column vector is
   $$\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$$

They are *jointly Gaussian* if the joint probability density is

$$p(x_1, x_2) = p(x) = \frac{1}{2\pi} \frac{1}{\sqrt{C}} e^{-\frac{1}{2}(x-\mu)^t H(x-\mu)} \ ,$$

where $H = C^{-1}$ is the *precision* matrix.

Random variables X and Y (Gaussian or not) are *independent* if $p_{XY}(x, y) = p_X(x)p_Y(y)$. Here, p_{XY} is the joint density, P_X is the density of X, and p_Y is the density of Y. Show that two jointly Gaussian random variables are uncorrelated if and only if they are independent.

4. Suppose X_1 and X_2 are correlated assets. Suppose we want to hedge X_1 with X_2. We use wX_2 as an approximation or hedge of X_1.

   (a) Calculate the optimal hedge, using variance to measure the quality of the hedge
   $$\min_w \text{var}(X_1 - wX_2) \ .$$
   Find a formula for the optimal w_* (the optimal w) in terms of variances and covariances.

   (b) Suppose the effectiveness of the hedge is measured by the variance reduction
   $$F = \frac{\text{var}(X_1 - w_* X_2)}{\text{var}(X_1)} \ .$$
   Find a formula for F as a function of ρ_{12}, the correlation coefficient between X_1 and X_2.

   (c) Use the result of part (b) to give a different proof of the Cauchy Schwarz inequality
   $$\text{cov}(X_1, X_2)^2 \leq \text{var}(X_1)\text{var}(X_2) \ .$$

5. Suppose there are two uncorrelated assets with $\text{var}(X_1) = \sigma_1^2$ and $\text{var}(X_2) = \sigma_2^2$ with the same expected return $\text{E}[X_1] = \text{E}[X_2] = \mu$. Suppose there is a risk free asset. Determine the efficient allocation between X_1 and X_2 (what fraction of the risky investment goes to X_1 and what fraction goes to X_2) as a function of σ_1 and σ_2. Why does the efficient allocation put some weight on an asset with the same return but more risk?

6. Suppose the covariance matrix and expected returns are
   $$C = .01 \cdot \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix} \ , \quad \mu = \begin{pmatrix} .1 \\ .2 \end{pmatrix} \ .$$

   (a) Find the minimum variance portfolio and the minimum variance. How much less is it than $\text{var}(X_1) = .01$?

   (b) What is the greatest expected return you can achieve with the risk of X_1?

(c) What is the smallest risk (variance) you can achieve with the expected return of X_1?

7. Suppose the minimum risk portfolio for a given expected return calls for a negative "position" in X_n. That is, $w_n < 0$. Some assets are easy to buy but hard to sell short (borrow). It it true that in this case, the optimal position is $w_n = 0$? That is, is it true that $w_n = 0$ has less variance for a given expected return than any portfolio with the same expected return and $w_n > 0$?

Computing exercise.

This exercise uses time series of historical stock prices to estimate the expected returns and covariance matrix. Then these are used to suggest efficient asset allocations using mean/variance analysis.

The historical data will come from *Alpha Vantage*. There is a link to the appropriate page of their site just below this assignment. The first step is to register with them and get an *API key*. This is free and quick, but you do have to give an email address. Click on *Claim your free API key*, do what it says, click *enter* and the new page will have the key, which is long string of capital letters and digits. Save this.

The second step is to download data using the key and an `R` script `PriceDownload.R` that is posted with this assignment. Near the top of that script is a place for you to copy in your personal API key. Then there is a list of *tickers*. A *ticker* is a short name of a company with a publicly traded stock, or some other public price. For example, the ticker for the Ford Motor Company is "F", and the ticker for an exchange traded ETF representing the S&P 500 index is "SPY". The script `PriceDownload.R` downloads a time series daily closing prices going back T days for each ticker. These time series are then stored in the same directory as the `PriceDownload.R` script in .csv files with names like `StockPriceTimeSeries_BA.csv` (this is the one for Boeing, with ticker "BA"). A *.csv* file (for *comma separated values*) is a simple format for spreadsheets. If you open one of these on your computer, you will get a spreadsheet program (Microsoft Excel, Apple Numbers, or something else). A .csv file is "human readable". If you open it with a text editor (xcode on an Apple laptop, xedit on a Linux box, dunno what on a Windows box), you will see column headings and numbers separated by commas. This is a common way to pass around financial data between computers with different operating systems (Linux, Windows, etc.).

The `PriceDownload.R` script script is more complicated than you might expect because the web site *Alpha Vantage* restricts you to five tickers per minute. Therefore, the script waits 70 seconds between every five downloads. That's the reason the download and analysis scripts are different. Downloading time consuming, so you don't want to do it over and over as you debug your analysis script.

The script *AssetAllocation.R* reads the .csv files and extracts the time series into an `R` matrix. It calculates the daily returns and, for each asset, the average daily return over the averaging period. This is `T = 250` days in the posted version. This is the *empirical* mean daily return. It is called μ_k here, and something similar in the script.

1. The next step is to compute the empirical covariance matrix for the same period. First subtract the empirical means to replace $R_{t,k}$ (the actual daily return for asset k at time t) by $R_{t,k} - \mu_k$. The empirical variances are

$$\sigma_{jj} = \frac{1}{T-2} \sum_{t=1}^{T-1} \left(R_{j,t} - \mu_j\right)^2 \ .$$

   If you have T prices, then you have $T-1$ daily returns and $T-2$ *degrees of freedom* in the variance estimate. Ask someone who has taken statistics why we use that $(T-2)$ instead of the number of data points $(T-1)$. It doesn't actually matter, let's not to upset people who believe in degrees of freedom. The empirical covariances are

$$\sigma_{jk} = \frac{1}{T-2} \sum_{t=1}^{T-1} \left(R_{j,t} - \mu_j\right)\left(R_{k,t} - \mu_k\right) \ .$$

   Assemble these into an $n \times n$ symmetric matrix C.

2. Add code to find efficient portfolios using mean/variance analysis. Find the minimum variance portfolio and the efficient portfolio that has the same expected return as the equal weighted portfolio. See how much the variance of the efficient portfolio differs from the variance of the equal weighted portfolio. Comment on the assets with negative weights.

3. Construct a ticker set of your choice by adding or removing tickers from the list in the posted version. Be sure that the ticker list in `AssetAllocation.R` is what you want. This determines the tickers that are used in analysis. You can use any asset that has a ticker. Possibilities incluce indices such as the S&P500, major international indices (China, Europe, developing, sectors, bond ETFs, REITs, etc.). They should be assets of your choosing that represent your interests. State your interest or interests and the say how the tickers you chose represent that interest or interests. Compare the mean/variance tradeoff of your list to the one in the posted script.

4. Compare the results with `250 days` of data to results with more or fewer days. How different are the portfolios, the covariance matrices, the expected returns, the variance of the efficient portfolio? For this, you have only to change the `T` variable in the `AssetAllocation.R` script. Be careful not to make T bigger than the actual time series. For example, don't ask for ten years of data for an asset that has existed only for five years.

5. Experiment to see the effect of putting in your own *views* of the expected returns. For example, you might think the expected return of the energy sector will be much less or more than the historical mean return. How

much does changing one expected return effect the portfolio allocation? Is the situation better or worse with longer datasets and more assets?

6. Do something else original and comment on the results.