

PORTFOLIO THEORY – LECTURE NOTES 1

Dr. Andrea Rigamonti

MEAN-VARIANCE OPTIMIZATION

From an economist's point of view, an investor that optimizes a portfolio is trying to maximize a utility function. An extremely simple utility function is the **linear utility function**:

$$U(V) = a + bV, \quad b > 0$$


where $U(V)$ is the utility that the investor gets depending on the value V of the portfolio. This function simply says that the higher the wealth, the higher the utility of the investor. Its shape is a line, and the parameter b determines how much the utility increases following a wealth increase. b is assumed to be positive, otherwise the investor would be indifferent ($b = 0$) or less satisfied ($b < 0$) when wealth increases. In other words, only the mean return of the portfolio matters.

Markowitz (1952) revolutionized the field by adding risk to the equation. His standard approach assumes that an investor cares not only about the mean but also the variance of the portfolio returns, i.e. the investor has a **mean-variance utility**. Given a certain mean return, the utility increases as the variance (which quantifies risk) gets lower. Equivalently, given a certain level of variance, the utility increases with a higher mean return. The decision in the trade-off between return and risk is quantified by a **risk aversion parameter** γ . A higher γ means that the investor is more risk averse and will therefore require a higher compensation for an increased risk. A lower γ means the investor has a lower risk-aversion and will be willing to take more risks. In other words, given a set of assets with a certain mean and variance, the lower the γ of the investor, the more he will create an optimal portfolio with a higher mean return but also a higher variance.

To model such preferences, a **quadratic utility function** is used:

$$U(V) = V - \frac{\gamma}{2}V^2, \quad \gamma > 0$$

γ is assumed to be positive so that the utility function is concave:



This implies that the investor is risk-averse. With $\gamma = 0$ the investor would be indifferent to risk, while $\gamma < 0$ would mean that the investor is risk taker (i.e., prefers more risk for the same amount of wealth, which is obviously not realistic).

Remember that the expected return of a portfolio is given by $\mu_p = \mathbf{w}'\boldsymbol{\mu} = V$, where \mathbf{w} is the vector of portfolio weights and $\boldsymbol{\mu}$ is the vector of mean return of the single assets. Moreover, recall that the variance is the expected value of the squared deviation from the mean.

So, the utility function becomes:

$$U(\mathbf{w}) = \mathbf{w}'\boldsymbol{\mu} - \frac{\gamma}{2}\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}$$

Therefore, given a risk-free asset and a set of N risky assets with mean returns $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, and a certain risk aversion coefficient γ , the investor we are considering wants to select the weights \boldsymbol{w} in a way that maximizes the following utility function:

$$\max_{\boldsymbol{w}} \boldsymbol{w}'\boldsymbol{\mu} - \frac{\gamma}{2} \boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}$$

This is an unconstrained optimization problem easy to solve. We just need to set the first-order condition, i.e., take the partial derivative with respect to \boldsymbol{w} and set it equal to zero:

$$\frac{\partial U(\boldsymbol{w})}{\partial \boldsymbol{w}} = \boldsymbol{\mu} - \frac{2\gamma}{2} \boldsymbol{\Sigma}\boldsymbol{w} = \boldsymbol{\mu} - \gamma\boldsymbol{\Sigma}\boldsymbol{w} = \mathbf{0}$$

We then solve for \boldsymbol{w} :

$$\begin{aligned} \boldsymbol{\Sigma}\boldsymbol{w} &= \frac{1}{\gamma} \boldsymbol{\mu} \\ \boldsymbol{w} &= \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} \end{aligned}$$

So, the closed-form solution that gives the optimal weights for the risky assets is:

$$\boldsymbol{w}_{mv} = \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$$

while the weight for the risk-free asset is equal to $1 - \boldsymbol{w}_{mv}'\mathbf{1}$, where $\mathbf{1}$ is a vector of ones with length equal to the number of risky assets.

The resulting optimal expected utility is:

$$U(\boldsymbol{w}_{mv}) = \frac{1}{2\gamma} \boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$$

Providing a specific value for γ might be difficult in practice. Moreover, it can be difficult to interpret the meaning of the specific utility value associated with a certain portfolio. While we can say that a portfolio with a certain utility is preferable to another with a lower utility, it is not obvious how good it is in a more general sense. In short, the value itself, in the case of utility, is not very informative.

A more intuitive approach is to specify the preferences through a desired mean portfolio return Re instead of a level of risk aversion. In this case, the goal becomes to **minimize the variance given the desired mean return**. This is a constrained optimization problem:

$$\min_{\boldsymbol{w}} \boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}$$

subject to:

$$\boldsymbol{w}'\boldsymbol{\mu} + (1 - \boldsymbol{w}'\mathbf{1})R_f = Re$$

In the constraint, $\boldsymbol{w}'\boldsymbol{\mu}$ is the return of the risky assets, and $(1 - \boldsymbol{w}'\mathbf{1})R_f$ is the return of the risk-free asset. Together they give the return of the portfolio, which, as stated, must be equal to Re .

To solve this problem, we use the method of Lagrange multipliers.

First we need to define the Lagrangian function, i.e., a modified version of the objective function that incorporates the constraint in this way:

$$L(\boldsymbol{w}, \lambda) = \boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w} + \lambda [Re - \boldsymbol{w}'\boldsymbol{\mu} - (1 - \boldsymbol{w}'\mathbf{1})R_f]$$

where λ is the Lagrange multiplier.

By including this additional term we can now solve an unconstrained problem instead of a constrained one. Therefore, we set the first order conditions for the Lagrangian function. The conditions involve two simultaneous equations, as we have to compute the partial derivative both with respect to \mathbf{w} and to λ .

$$\frac{\partial L}{\partial \mathbf{w}} = 2\Sigma\mathbf{w} - \lambda\boldsymbol{\mu} + \lambda R_f \mathbf{1} = \mathbf{0}$$

$$\frac{\partial L}{\partial \lambda} = Re - \mathbf{w}'\boldsymbol{\mu} - (1 - \mathbf{w}'\mathbf{1})R_f = 0$$

Solving this system for \mathbf{w} is somewhat computation intensive, but it gives a closed-form solution. Specifically, the optimal portfolio weights in this case are:

$$\mathbf{w}_{mv} = \frac{Re}{\boldsymbol{\mu}'\Sigma^{-1}\boldsymbol{\mu}}\Sigma^{-1}\boldsymbol{\mu}$$

In both cases we do not need to explicitly include the risk-free asset in the asset menu, as it is equivalent and simpler to work with excess returns, i.e. with the returns of the risky assets from which we subtracted the risk-free rate.

Obviously, it is also possible to specify a given level of variance and maximize the mean return. Mathematically, these are two equivalent optimization problems. However, it is more intuitive and more common to specify the desired mean and minimize the variance.

So far we treated the inputs $\boldsymbol{\mu}$ and Σ as if they are given, but in practice they need to be estimated. The simplest approach is called **plug-in approach**: the sample estimates of the inputs are computed from past data, and are then plugged into the optimization problem as if they were the true values. Obviously, this is not really the case, as sample estimates can be poor estimates of the true parameter values. In particular, the estimation error in the sample mean is typically so big that minimizing the variance while ignoring $\boldsymbol{\mu}$ usually leads to portfolios with a higher Sharpe ratio than those computed via mean-variance optimization.

Therefore, the investor might want to compute the **global minimum variance portfolio (GMV)**, also simply called **minimum variance portfolio**. Obviously, we need to impose the constraint that the sum of the weights of the risky assets is equal to one, which means that nothing is invested in the risk-free asset. Otherwise, everything would be invested in the risk-free asset. Hence, we have to solve the following constrained optimization problem:

$$\min_{\mathbf{w}} \mathbf{w}'\Sigma\mathbf{w}$$

subject to:

$$\mathbf{w}'\mathbf{1} = 1$$

We write the Lagrangian function:

$$L(\mathbf{w}, \lambda) = \mathbf{w}'\Sigma\mathbf{w} + \lambda[1 - \mathbf{w}'\mathbf{1}]$$

In order to get a more convenient first order condition, it is common to multiply the first term by 0.5, which does not alter the result. So the Lagrangian becomes:

$$L(\mathbf{w}, \lambda) = \frac{1}{2}\mathbf{w}'\Sigma\mathbf{w} + \lambda[1 - \mathbf{w}'\mathbf{1}]$$

The first order conditions are:

$$\frac{\partial L}{\partial \mathbf{w}} = \Sigma\mathbf{w} - \lambda\mathbf{1} = \mathbf{0}$$

$$\frac{\partial L}{\partial \lambda} = 1 - \mathbf{w}'\mathbf{1} = 0$$

Through some simple rearrangement we get:

$$\begin{aligned}\mathbf{w} &= \lambda \Sigma^{-1} \mathbf{1} \\ \mathbf{w}'\mathbf{1} &= 1\end{aligned}$$

In the first equation, we can multiply both sides by $\mathbf{1}'$, obtaining:

$$\mathbf{1}'\mathbf{w} = \lambda \mathbf{1}'\Sigma^{-1}\mathbf{1}$$

From the second equation we know that:

$$\mathbf{w}'\mathbf{1} = \mathbf{1}'\mathbf{w} = 1$$

Hence, the first equation becomes:

$$\begin{aligned}1 &= \lambda \mathbf{1}'\Sigma^{-1}\mathbf{1} \\ \lambda &= \frac{1}{\mathbf{1}'\Sigma^{-1}\mathbf{1}}\end{aligned}$$

So, finally, we can take this last result and replace λ in $\mathbf{w} = \lambda \Sigma^{-1} \mathbf{1}$, obtaining:

$$\mathbf{w} = \frac{1}{\mathbf{1}'\Sigma^{-1}\mathbf{1}} \Sigma^{-1} \mathbf{1}$$

Therefore, the closed form-solution that gives the minimum variance weights is:

$$\mathbf{w}_v = \frac{1}{\mathbf{1}'\Sigma^{-1}\mathbf{1}} \Sigma^{-1} \mathbf{1}$$

As nothing is invested in the risk-free rate (since we require that the weights for the risky assets must sum up to 1), it is equivalent to work with returns or excess returns. However, it might be convenient to still work with excess returns, so that the results will be easily comparable with those obtained by the mean-variance portfolio.

A further improvement can come from restricting the minimum variance portfolio to only have long positions. In other words, we add another constraint that prohibits short selling positions, to get a **long-only minimum variance portfolio**. This problem does not have a closed form solution, but it can easily be solved with computer programs.

When a value for γ is specified, another strategy that mitigates the impact of the estimation error is the **1/N rule**.¹ In this rule the (sample) estimates of $\boldsymbol{\mu}$ and Σ are used to optimally allocate the wealth between the risk-free asset and the equally weighted risky assets. This rule usually performs very well. The weights for the risky assets according to this strategy are:

$$\mathbf{w}_{1/N} = \frac{1}{\gamma} \frac{\mathbf{1}'\boldsymbol{\mu}}{\mathbf{1}'\Sigma\mathbf{1}} \mathbf{1}$$

and the weight for the riskless asset is given by $1 - \mathbf{w}'_{1/N}\mathbf{1}$.

For example, if $N = 5$ and this rule returns a weight of 0.15 for each risky asset, we equally divide 75% of our wealth among the risky assets (i.e., 15% on each risky asset), and then place the remaining 25% on the risk-free asset.

¹ The name "1/N rule" is often used to refer to a naive rule where one simply places all the wealth on risky assets with all weights equal to 1/N, without estimating any input. In these notes we refer instead to the more elaborate rule described in the text.