WILMOTT
TECHNICAL ARTICLE

PUBLISHED ON WILMOTT.COM IN OCTOBER 2001

SERVING THE QUANTITATIVE FINANCE COMMUNITY

Technical article for www.wilmott.com subscribers

A COMMON PITFALL IN MEAN-VARIANCE ASSET ALLOCATION

Attilio Meucci

Author details:

Attilio Meucci is head of quantitative research at Bain & Co, Milan office, Via Crocefisso 10 I-20212.

Email: attilio.meucci@bca.it

Abstract

The most popular approach to asset allocation is by far Markowitz's efficient frontier framework, where the investor's goal is to maximize a mean-variance utility function. Even though this problem has been thoroughly studied and implemented worldwide, it is still common to fall into misunderstandings due to an inappropriate use of the definitions of returns on assets: these misunderstandings lead to sub-optimal asset allocations. In this paper, first we define the proper theoretical framework, then we compare in an example the asset allocation obtained in the correct and the incorrect approach respectively. The differences are significant, especially for long investment horizons. Furthermore, we obtain the surprising result that, under standard hypotheses, a long-horizon allocation should invest in less volatile assets than a short-horizon allocation, contrary to the common belief that "for the long run stocks are a better investment than bonds"'.

A common pitfall in mean-variance asset allocation

Attilio Meucci*

The most popular approach to asset allocation is by far Markowitz's efficient frontier framework, where the investor's goal is to maximize a mean-variance utility function. Even though this problem has been thoroughly studied and implemented worldwide, it is still common to fall into misunderstandings due to an inappropriate use of the definitions of returns on assets: these misunderstandings lead to sub-optimal asset allocations. In this paper, first we define the proper theoretical framework, then we compare in an example the asset allocation obtained in the correct and the incorrect approach respectively. The differences are significant, especially for long investment horizons. Furthermore, we obtain the surprising result that, under standard hypotheses, a long-horizon allocation should invest in less volatile assets than a short-horizon allocation, contrary to the common belief that "for the long run stocks are a better investment than bonds".

1 Introduction

Markowitz's mean-variance approach (Markowitz [1959]) can be summarized as follows: find the combination of the given assets that maximizes the expected return on a portfolio π at a specified time horizon τ for a given level of risk aversion k. This optimization problem can be written as

$$\max_{\substack{\text{feasible} \\ \text{combinations of assets}}} \left(E_t \left(R^{\pi}_{\tau} \right) - k \cdot var_t \left(R^{\pi}_{\tau} \right) \right), \tag{1}$$

where R^{π} is the portfolio return at the given horizon, the constant k is the risk aversion $E_t(\cdot)$ is the expected value at the current date t and $var_t(\cdot)$ is the expected variance. Before we can solve this problem, we need a precise definition of returns: we have to distinguish between *logarithmic* (or compounded) and *linear* returns. Consider a generic asset whose price we denote P_t. *Logarithmic* returns $r_{t,\tau}$ at time t with horizon τ are defined in such a way that $P_{t+\tau} = e^{r_{t,\tau}} P_t$, therefore

$$r_{t,\tau} = \ln(\frac{P_{t+\tau}}{P_t}). \tag{2}$$

On the other hand *linear* returns $R_{t,\tau}$ at time t with horizon τ are defined in such a way that $P_{t+\tau} = (1 + R_{t,\tau}) P_t$, therefore

$$R_{t,\tau} = \frac{P_{t+\tau}}{P_t} - 1. \tag{3}$$

*Attilio Meucci, Ph.D., is head of quantitative research at Bain & Co, Milan office, via Crocefisso 10 I-20212. attilio.meucci@bca.it

The classical mean-variance optimization problem (1) is set in terms of *linear* returns. Of course, one could as well define it in terms of logarithmic returns, but two major problems would arise: the links with utility theory would be lost (Ingersoll [1987]), and the mathematics would be much harder. Indeed, if we denote by ω_i the i-th asset 's relative weight in our portfolio it is straightforward to check that the *linear* return on the portfolio at the horizon τ is the weighted average of the *linear* returns $\mathbf{R}_\tau$ on the assets: $R_\tau^\pi = \omega'\mathbf{R}_\tau$. Therefore the expected value of the *linear* return on the portfolio can be expressed in terms of the expected values $\mathbf{M}_\tau$ of the *linear* returns on assets as $E_t(R_\tau^\pi) = \omega'\mathbf{M}_\tau$, and the expected variance of the *linear* return on the portfolio can be expressed in terms of the expected covariance matrix S_τ of the *linear* returns on assets as $var_t(R_\tau^\pi) = \omega' S_\tau \omega$. This is not true if we set problem (1) in terms of logarithmic returns. If the horizon is relatively short, say of the order of one year, the differences that arise from misspecifying the returns are negligible. The situation changes at longer horizon, as we show here.

2 The correct framework

In this section we go through all the steps of an asset allocation problem as faced by a practitioner, highlighting the mistakes that arise from misspecifying the returns and fixing them on the way.

Step 1: We choose a set of assets to form our portfolio (in our example, US bonds 1-3yrs, US bonds 7-10yrs, US stocks utilities, US stocks financial, US stocks technology), we set an investment horizon τ and a level of risk aversion.

Step 2: We input the expected covariance matrix of *linear* returns on assets at the specified horizon S_τ. One option is to estimate it directly based on time series analysis, but we typically have too few observations when the horizon is large. Therefore we need to derive it by means of models, which are usually set in terms of *logarithmic* returns. In our example we assume that asset prices follow a lognormal stochastic process. Therefore, first we estimate from our ten year time series the covariance matrix of *logarithmic* monthly returns on assets. $\Sigma_{1/12}$ by means of the standard formula; then we model the covariance matrix at the desired horizon $\Sigma_\tau = \Sigma_{1/12} \cdot 12 \cdot \tau$. Of course, other models could be assumed to obtain Σ_τ (e.g., weak mean reversion of returns, interpolation of weighted estimates at different horizons,...).

Step 3: We input the expected *linear* returns. In common practice, one expresses views on returns by means of some model in terms of one-year horizons, which is tantamount to inputting expected one-year *linear* returns on assets $\mathbf{M}_1 = E(\mathbf{R}_{t,1})$. To illustrate, in our example we used a simple risk-premium model, that assumes these returns are one third of the yearly volatility.

Step 4 (Incorrect!): We compute the expected *linear* returns on assets at the horizon as $\mu_\tau = \mathbf{M}_1 \cdot \tau$ and solve (1) for the best portfolio:

$$\max_\omega (\omega'\mu_\tau - k\omega'\Sigma_\tau\omega) \qquad (4)$$
$$\text{s.t. } \sum \omega_i = 1$$

We stress that (4) is incorrect: Σ_τ is not the covariance matrix of linear returns and μ_τ are not the expected linear returns. Furthermore, μ_τ are not even expected logarithmic returns (only a first order approximation). But even if

μ_τ were expected logarithmic returns, $\omega' \mu_\tau$ is not the expected value of portfolio's logarithmic return and $\omega' \Sigma_\tau \omega$ is not the expected variance of portfolio's logarithmic return:

$$\begin{aligned} E_t\left(r^{\pi}_{t,\tau}\right) &\neq \omega' \mu_\tau \\ var_t\left(r^{\pi}_{t,\tau}\right) &\neq \omega' \Sigma_\tau \omega \end{aligned} \tag{5}$$

Therefore(4) would not solve (1) even if it were stated in terms of logarithmic returns.

Step 4 (Correct!): Therefore, instead of solving (4), we need to compute from our inputs $\mathbf{M}_1$ and Σ_τ the expected values and covariance matrix of *linear* returns at the given horizon $\mathbf{M}_\tau$and S_τ(see the appendix) and then solve problem (1) the proper way:

$$\begin{aligned} &\max_{\omega} \left(\omega' \mathbf{M}_\tau - k \omega' S_\tau \omega\right) \\ &\text{s.t.} \ \sum \omega_i = 1 \end{aligned} \tag{6}$$

In the figures we compare the correct and the incorrect approach.

In Figure 1 the investment horizon is one year. In the plot on the left we show the correct and the incorrect efficient frontiers, i.e., the solutions to the correct (6) and the incorrect (4) maximization problems for all levels of risk aversion k. On each frontier we select a specific portfolio, which corresponds to a given value of risk aversion ($k = 2$). In the plot on the right, we show the asset allocation corresponding to that portfolio in the correct and the incorrect approach respectively. Since the investment horizon is relatively short, i.e., one year, logarithmic and linear returns are approximately the same and the differences in terms of both efficient frontiers and asset allocations are negligible.

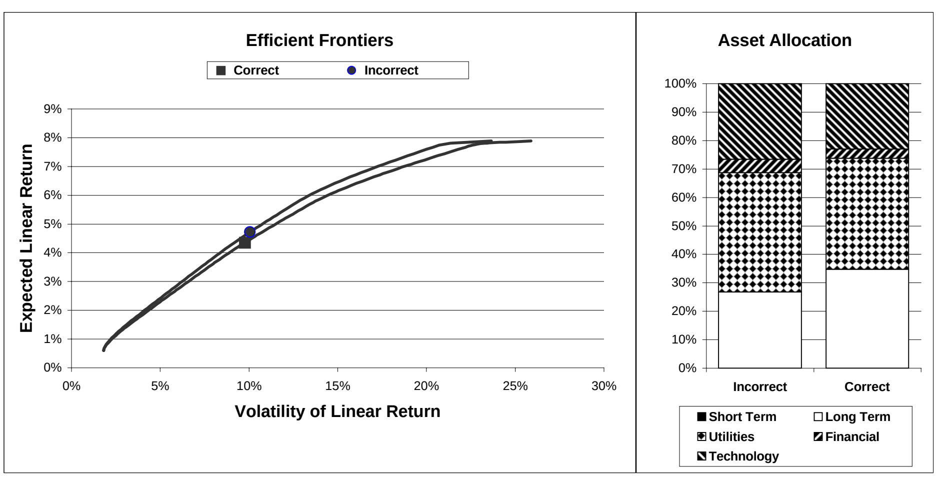

One-year horizon.

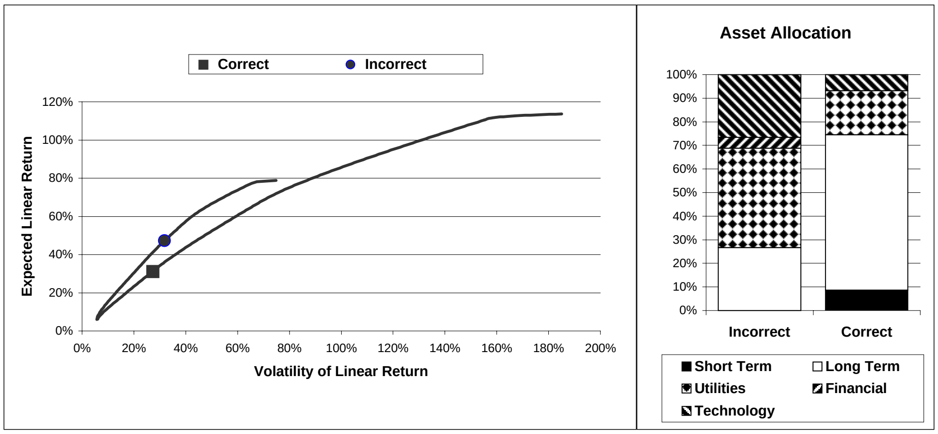

Ten-year horizon.

As the investment horizon becomes larger, the difference between linear and logarithmic returns gives rise to completely different scenarios. In Figure 2 we shift the investment horizon from one to ten years and then we plot efficient frontiers and optimal asset allocation of a specific portfolio just as we did in Figure 1. The incorrect efficient frontier, plotted as solutions to the incorrect maximization problem (4) for all levels of risk aversion k, is now completely distorted with respect to the correct efficient frontier that solves (6). As for the optimal asset allocation, the incorrect approach yields exactly the same asset allocation as it did at a short horizon, due to the horizon invariance of (4) under the lognormal hypothesis. The correct approach, on the other hand, overweighs less volatile assets, as opposed to the common belief that "for the long run stocks are a better investment than bonds".

3 Conclusions

Markowitz's mean-variance approach to optimal asset allocation seeks to maximize a portfolio's expected *linear* return minimizing the volatility of the portfolio's *linear* return. Practitioners instead typically use *logarithmic* returns: this is mathematically inconsistent and the resulting asset allocation is suboptimal. To illustrate, we show the optimal asset allocation resulting from the standard assumption that the processes followed by the asset prices are lognormal: solving the problem the incorrect way, i.e., with logarithmic returns, we obtain a horizon-invariant asset allocation; solving the problem the correct way, i.e., with linear returns, we obtain the surprising result that optimal long-horizon investments are composed by less volatile assets than short-horizon investments.

Acknowledgement

The author would like to thank Gianluca Fusai for his insightful comments on the first draft of the paper.

A Appendix

Here we provide a way to compute the covariance matrix and expected values of linear returns for a generic distribution of the logarithmic returns $\psi(\mathbf{r})$. Our aim is computing the expected values and covariance matrix of linear returns at all horizons τ

$$\mathbf{M}_\tau = E(\mathbf{R}_{t,\tau}), \qquad S_\tau = cov(\mathbf{R}_{t,\tau}).$$

First we express the linear returns in terms of the logarithmic ones: from (2) and (3) we see that $R_{t,\tau} = e^{r_{t,\tau}} - 1$. Therefore

$$M_\tau^i = E\left(R_{t,\tau}^i\right) = \int R_{t,\tau}^i(\mathbf{r})\,\psi(\mathbf{r})\,d\mathbf{r} = \int \left(e^{r_{t,\tau}^i} - 1\right)\psi(\mathbf{r})\,d\mathbf{r}$$

Similarly

$$\begin{aligned} E\left(R_{t,\tau}^i R_{t,\tau}^j\right) &= \int R_{t,\tau}^i(\mathbf{r})\,R_{t,\tau}^j(\mathbf{r})\,\psi(\mathbf{r})\,d\mathbf{r} = \int \left(e^{r_{t,\tau}^i} - 1\right)\left(e^{r_{t,\tau}^j} - 1\right)\psi(\mathbf{r})\,d\mathbf{r} \\ &= \int e^{r_{t,\tau}^i + r_{t,\tau}^j}\psi(\mathbf{r})\,d\mathbf{r} - \int e^{r_{t,\tau}^i}\psi(\mathbf{r})\,d\mathbf{r} - \int e^{r_{t,\tau}^j}\psi(\mathbf{r})\,d\mathbf{r} + 1 \end{aligned}$$

We notice that the above quantities can be easily expressed in terms of the moment-generating function $\phi(\mathbf{t}) = \int e^{\mathbf{t}'\mathbf{r}}\psi(\mathbf{r})\,d\mathbf{r}$ of the distribution ψ. In our lognormal economy the moment-generating function reads $\phi_{\mu,\Sigma}(\mathbf{t}) = e^{\mathbf{t}'\mu + \frac{1}{2}\mathbf{t}'\Sigma\mathbf{t}}$. Therefore, the expected linear returns turn out to be

$$M_\tau^i = e^{\mu_\tau^i + \frac{1}{2}\Sigma_\tau^{ii}} - 1 \tag{7}$$

As for the second moment, we obtain

$$E\left(R_{t,\tau}^i R_{t,\tau}^j\right) = e^{\mu_\tau^i + \mu_\tau^j + \frac{1}{2}\left(\Sigma_\tau^{ii} + \Sigma_\tau^{jj} + 2\Sigma_\tau^{ij}\right)} - e^{\mu_\tau^i + \frac{1}{2}\Sigma_\tau^{ii}} - e^{\mu_\tau^j + \frac{1}{2}\Sigma_\tau^{jj}} + 1,$$

and thus

$$S^{ij} = E\left(R_{t,\tau}^i R_{t,\tau}^j\right) - M_\tau^i M_\tau^j = e^{\mu_\tau^i + \mu_\tau^j + \frac{1}{2}\left(\Sigma_\tau^{ii} + \Sigma_\tau^{jj}\right)}\left(e^{\Sigma_\tau^{ij}} - 1\right) \tag{8}$$

Equations (7) and (8) represent the correct parameters for the maximization problem (6), in terms of the covariance matrix of logarithmic returns Σ_τ and the expected logarithmic returns $\mu_\tau = E(\mathbf{r}_{t,\tau})$. Since our original inputs are Σ_τ and the one-year expected linear returns $\mathbf{M}_1 = E(\mathbf{R}_{t,1})$, we still have to derive an expression for μ_τ. Again, we assume that asset prices follow a lognormal stochastic process, but other models could be considered. Under this hypothesis there is a ν such that $e^\nu = \mathbf{1} + \mathbf{M}_1$ and $E(\mathbf{r}_{t,1}) = \nu - diag(\Sigma_1)/2$. Eliminating ν we obtain

$$\mu_\tau = \tau\mu_1 = \tau\ln(\mathbf{1} + \mathbf{M}_1) - \frac{\tau}{2}diag(\Sigma_1). \tag{9}$$

Inserting (9) in (7) and (8) completely determines the parameters to solve the optimization (6).

References

Ingersoll, J.E. Jr. *Theory of Financial Decision Making*. Rowman & Littlefield, 1987.

Markowitz, H. *Portfolio Selection: efficient diversification of investment*. John Wiley, 1959.

SERVING THE QUANTITATIVE FINANCE COMMUNITY

Phone: 44 (0) 20 7792 1310
Fax: 44 (0) 7050 670002
Email: tech@wilmott.com

Wilmott website and newsletter are regularly read by thousands of quants from all the corners of the globe. Readers can be found in investment banks, hedge funds, consultancies, software companies, pension funds and academia.

Our community of finance professionals thrives on cutting-edge research, innovative models and exciting new products. Their thirst for knowledge is almost unquenchable.

To keep our readers happy we are proactive in seeking out the best new research and the brightest new researchers. However, even our eagle eyes and relentless searching occasionally miss something of note.

If you are active in quantitative finance and would like to submit work for publication on the site, or for mention in the newsletter, contact us at **submit@wilmott.com**.

VISIT WWW.WILMOTT.COM FOR

- TECHNICAL ARTICLES—CUTTING-EDGE RESEARCH, IMPLEMENTATION OF KEY MODELS
- LYCEUM—EDUCATIONAL MATERIAL TO HELP YOU GET UP TO SPEED
- BOOKSHOP—GREAT OFFERS ON ALL THE IMPORTANT QUANTITATIVE FINANCE BOOKS
- FORUM—MOST ACTIVE QUANTITATIVE FINANCE FORUM ON THE INTERNET

WILMOTT
SERVING THE QUANTITATIVE FINANCE COMMUNITY