Portfolio Choice (1)

John Y. Campbell

Ec2723

September 2010

The Principle of Participation

Consider the choice between one safe and one risky asset. An investor with initial wealth w can invest in a safe asset with return r, or a risky asset with return $r + \widetilde{x}$. Final wealth is

$$w(1 + r) + \theta\widetilde{x} = w_0 + \theta\widetilde{x},$$

where θ is the dollar amount (not the share of wealth) invested in the risky asset. The investor's problem is

$$\mathrm{Max}_\theta \ V(\theta) = \mathrm{E}u(w_0 + \theta\widetilde{x}).$$

The Principle of Participation

The first-order condition is

$$V'(\theta^*) = \mathrm{E}\widetilde{x}u'(w_0 + \theta^*\widetilde{x}) = 0.$$

We have

$$V'(0) = \mathrm{E}\widetilde{x}u'(w_0),$$

which has the same sign as $\mathrm{E}\widetilde{x}$. The investment in the risky asset should be positive if it has a positive expected return. This is true for any level of risk aversion. Thus we cannot explain non-participation in risky asset markets by risk aversion. We need fixed costs of participation or a kink in the utility function that generates "first-order risk aversion".

Portfolio Choice with a Small Risk

We consider a small risk

$$\widetilde{x} = k\mu + \widetilde{y}$$

and assume $k > 0$. The first-order condition is

$$\mathrm{E}(k\mu + \widetilde{y})u'(w_0 + \theta^*(k)(k\mu + \widetilde{y})) = 0.$$

Differentiating w.r.t. k,

$$\mu \mathrm{E}u'(\widetilde{w}) + \theta^*(k)\mu \mathrm{E}(k\mu + \widetilde{y})u''(\widetilde{w}) + \theta^{*\prime}(k)\mathrm{E}(k\mu + \widetilde{y})^2 u''(\widetilde{w}) = 0.$$

Evaluating at $k = 0$,

$$\theta^{*\prime}(0) = \frac{\mu}{\mathrm{E}\widetilde{y}^2}\frac{1}{A(w_0)}.$$

Portfolio Choice with a Small Risk

Then a Taylor expansion for the investment in the risky asset gives

$$\theta^*(k) \approx \theta^*(0) + k\theta^{*\prime}(0) = \frac{\mathrm{E}\widetilde{x}}{\mathrm{E}(\widetilde{x} - \mathrm{E}\widetilde{x})^2} \frac{1}{A(w_0)}.$$

We can divide θ by wealth to find the share of wealth invested in the risky asset. Call this α. We find

$$\alpha^*(k) = \frac{\theta^*(k)}{w_0} \approx \frac{\mathrm{E}\widetilde{x}}{\mathrm{E}(\widetilde{x} - \mathrm{E}\widetilde{x})^2} \frac{1}{R(w_0)}.$$

Portfolio Choice in the CARA Normal Case

The above formula for dollars invested in the risky asset is exact when risk is normal, $\widetilde{x} \sim N(\mu, \sigma^2)$, and utility is CARA with risk aversion A. In this case the problem becomes

$$\mathrm{Max}_\theta \ V(\theta) = \mathrm{E}[-\exp(-A(w_0 + \theta\widetilde{x}))].$$

Utility is lognormally distributed (its log is normally distributed) if $\widetilde{x}$ is normally distributed. For any lognormal random variable $\widetilde{z}$, we have the following extremely useful result:

$$\log \mathrm{E}(\widetilde{z}) = \mathrm{E}\log(\widetilde{z}) + \frac{1}{2}\mathrm{Var}\log(\widetilde{z}).$$

Portfolio Choice in the CARA Normal Case

The portfolio choice problem is equivalent to

$$\text{Min} \log \mathrm{E}\left[\exp(-A(w_0 + \theta\widetilde{x}))\right] = -A(w_0 + \theta\mu) + \frac{1}{2}A^2\theta^2\sigma^2,$$

which is equivalent to

$$\text{Max } A(w_0 + \theta\mu) - \frac{1}{2}A^2\theta^2\sigma^2.$$

The solution is

$$\theta^* = \frac{\mu}{A\sigma^2},$$

independent of the initial level of wealth.

Portfolio Choice in the CARA Normal Case

This framework is very tractable:

- It is easy to add multiple assets.
- It is easy to handle additive background risk, arising from random income or nontradable assets. Background risk does not affect the demand for tradable risky assets if it is uncorrelated with their returns.
- Equilibrium with heterogeneous agents is easy to calculate because the wealth distribution does not affect the demand for risky assets.

Portfolio Choice in the CARA Normal Case

However there are also serious problems with this framework:

- Wealth does not affect the amount invested in a risky asset.
- Growth in consumption and wealth with multiplicative risks implies increasing absolute risks. CARA implies that this should generate an upward trend in risk premia which we have not seen historically.
- The assumption of normality cannot hold over more than one time interval. The compounding of returns over many periods converts a symmetric normal distribution into a right-skewed, non-normal distribution.

Portfolio Choice in the CRRA Lognormal Case

Assume that asset returns are lognormally distributed, and utility is power with relative risk aversion γ. The maximization problem is

$$\max \mathrm{E}_t W_{t+1}^{1-\gamma}/(1-\gamma).$$

When $\gamma < 1$, this expectation is positive and maximizing it is equivalent to maximizing its log. (When $\gamma > 1$, then we flip sign, minimize the log, and get the same solution.) Proceeding with the $\gamma < 1$ case, if next-period wealth is lognormal, we can rewrite the problem as

$$\max \log \mathrm{E}_t W_{t+1}^{1-\gamma} = (1-\gamma)\mathrm{E}_t w_{t+1} + \frac{1}{2}(1-\gamma)^2\sigma_{wt}^2 - \log(1-\gamma),$$

where $w_t = \log(W_t)$. This has the same solution as

$$\max \mathrm{E}_t w_{t+1} + \frac{1}{2}(1-\gamma)\sigma_{wt}^2.$$

Portfolio Choice in the CRRA Lognormal Case

The budget constraint is

$$w_{t+1} = r_{p,t+1} + w_t,$$

where $r_{p,t+1} = \log(1 + R_{p,t+1})$ is the log return on the portfolio. So we can restate the problem as

$$\max \mathrm{E}_t r_{p,t+1} + \frac{1}{2}(1-\gamma)\sigma_{pt}^2,$$

where σ_{pt}^2 is the conditional variance of the log portfolio return.

Portfolio Choice in the CRRA Lognormal Case

To understand this, note that

$$\mathrm{E}_t r_{p,t+1} + \sigma_{pt}^2/2 = \log \mathrm{E}_t(1 + R_{p,t+1}).$$

Thus we can equivalently write

$$\max \log \mathrm{E}_t(1 + R_{p,t+1}) - \frac{\gamma}{2}\sigma_{pt}^2.$$

The investor trades off the log of the *arithmetic* mean return against the variance of the log return, and the variance penalty is proportional to risk aversion.

Log Utility and the Growth-Optimal Portfolio

When $\gamma = 1$, the investor has log utility and chooses the *growth-optimal portfolio* with the maximum log return.

- When $\gamma > 1$, the investor seeks a safer portfolio by penalizing the variance of log returns.
- When $\gamma < 1$, the investor actually seeks a riskier portfolio because a higher variance, with the same mean log return, corresponds to a higher mean simple return.
- The case $\gamma = 1$ is the boundary where these two opposing considerations cancel.

An Attractive Property of the Growth-Optimal Portfolio

The growth-optimal portfolio has the property that as the investment horizon increases, it outperforms any other portfolio with increasing probability.

- The difference between the log return on the growth-optimal portfolio and the log return on any other portfolio is normally distributed with a positive mean.
- Assume that returns are iid over time.
- Then as the horizon increases, the mean and variance of the excess log return both grow linearly, so the ratio of mean to standard deviation grows with the square root of the horizon.
- The ratio of mean to standard deviation determines the probability that the excess return is positive, which therefore increases with the investment horizon.
- Markowitz and others have used this to argue that long-term investors should have log utility, but this claim has been strongly opposed by Samuelson and others.

Log Asset Returns and Log Portfolio Returns

Now we need to relate the log portfolio return to the log returns on underlying assets. The simple return on the portfolio is a linear combination of the simple returns on the risky and riskless assets. The log return on the portfolio is the log of this linear combination, which is not the same as a linear combination of logs.

Over short time intervals, however, we can use a Taylor approximation of the nonlinear function relating log individual-asset returns to log portfolio returns:

$$r_{p,t+1} - r_{f,t+1} \approx \alpha_t (r_{t+1} - r_{f,t+1}) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma_t^2.$$

The difference between the log portfolio return and a linear combination of log individual-asset returns is given by $\alpha_t(1-\alpha_t)\sigma_t^2/2$ which is zero if $\alpha_t = 0$ or 1. When $0 < \alpha_t < 1$, the portfolio is a weighted average of the individual assets and the term $\alpha_t(1-\alpha_t)\sigma_t^2/2$ is positive because the log of an average is greater than an average of logs.

Log Asset Returns and Log Portfolio Returns

Another way to understand this is to rewrite the equation as

$$r_{p,t+1} - r_{f,t+1} + \frac{\sigma_{pt}^2}{2} \approx \alpha_t \left(r_{t+1} - r_{f,t+1} + \frac{\sigma_t^2}{2} \right),$$

using the fact that $\sigma_{pt}^2 = \alpha_t^2 \sigma_t^2$. This shows that the mean of the simple excess portfolio return is linearly related to the mean of the simple excess return on the risky asset.

Properties of this approximation:

- It becomes more accurate as the time interval shrinks. It is exact in continuous time with continuous paths for asset prices (then it follows from Itô's Lemma.).
- It rules out bankruptcy, even with a short position ($\alpha_t < 0$) or leverage ($\alpha_t > 1$).

Back to the Portfolio Choice Problem

With two assets, the mean excess portfolio return is $E_t r_{p,t+1} - r_{f,t+1} = \alpha_t (E_t r_{t+1} - r_{f,t+1}) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma_t^2$, while the variance of the portfolio return is $\alpha_t^2\sigma_t^2$. Substituting into the objective function, we get

$$\max \alpha_t (E_t r_{t+1} - r_{f,t+1}) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma_t^2 + \frac{1}{2}(1-\gamma)\alpha_t^2\sigma_t^2.$$

The solution is

$$\alpha_t = \frac{E_t r_{t+1} - r_{f,t+1} + \sigma_t^2/2}{\gamma\sigma_t^2},$$

which is an exact version of the result for a small risk.

Mean-Variance Analysis

Mean-variance analysis judges portfolios by their first two moments of returns (a good starting point, but not the end of the story!)
In a static (single-period) model, this can be justified by

- Quadratic utility of wealth, or
- Return distributions for which the first two moments are sufficient statistics. E.g.
  - Normal distribution
  - Lognormal distribution (with short time intervals so that portfolio returns and individual asset returns can both be lognormal)
  - Multivariate t distribution
  - Any of the above, plus an arbitrarily distributed common risk that does not affect portfolio choice.

Mean-Variance Analysis

Along with assumptions on return distributions, we may also use utility assumptions (e.g. CARA-normal, CRRA-lognormal) to get closed-form portfolio rules. But these utility assumptions are not needed to justify mean-variance analysis.

For tractability, we assume that short sales are permitted.

Two Risky Assets

Start with the choice between two risky assets with returns R_1 and R_2.

$$R_p = w_1 R_1 + w_2 R_2.$$

$$\overline{R}_p = w_1 \overline{R}_1 + w_2 \overline{R}_2.$$

$$\begin{aligned}
\sigma_p^2 &= \mathrm{Var}(w_1 R_1 + w_2 R_2) \\
&= w_1^2 \mathrm{Var}(R_1) + w_2^2 \mathrm{Var}(R_2) + 2 w_1 w_2 \mathrm{Cov}(R_1, R_2) \\
&= w_1^2 \sigma_1^2 + w_2^2 \sigma_2^2 + 2 w_1 w_2 \sigma_{12} \\
&= w_1^2 \sigma_1^2 + w_2^2 \sigma_2^2 + 2 w_1 w_2 \sigma_1 \sigma_2 \rho_{12} \ ,
\end{aligned}$$

where $\rho_{12} \equiv \mathrm{Corr}(R_1, R_2)$.

Two Risky Assets

Given a target $\overline{R}_p$,

$$w_1 = \frac{\overline{R}_p - \overline{R}_2}{\overline{R}_1 - \overline{R}_2} \, .$$

The variance of the portfolio return is

$$\sigma_p^2 = w_1^2 \sigma_1^2 + (1 - w_1)^2 \sigma_2^2 + 2 w_1 (1 - w_1) \sigma_1 \sigma_2 \rho_{12} \, .$$

This is a quadratic function of w_1, and hence of $\overline{R}_p$.

$$\begin{aligned} \frac{d\sigma_p^2}{dw_1} &= 2 w_1 \sigma_1^2 - 2(1 - w_1)\sigma_2^2 + 2(1 - 2w_1)\sigma_{12} \\ &= 2 w_1 [\sigma_1^2 + \sigma_2^2 - 2\sigma_{12}] - 2[\sigma_2^2 - \sigma_{12}] \, . \end{aligned}$$

This derivative is increasing in w_1.

Risk-Return Plots

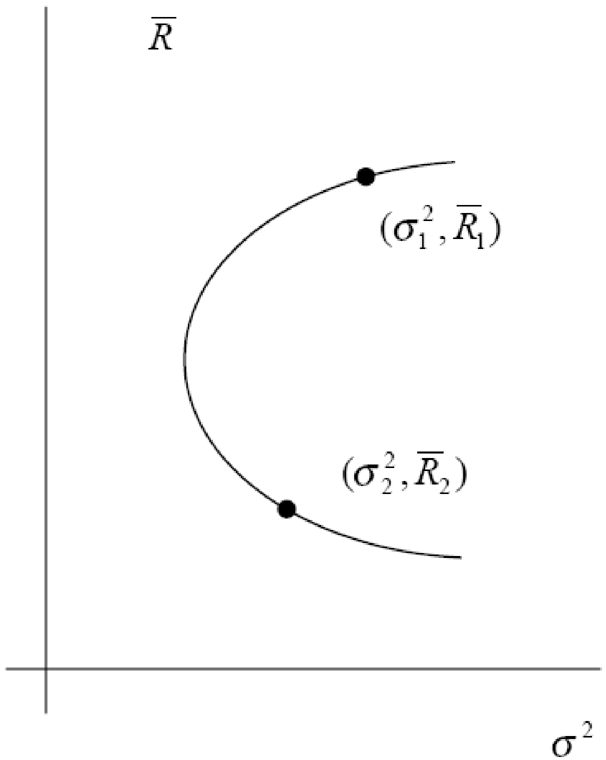

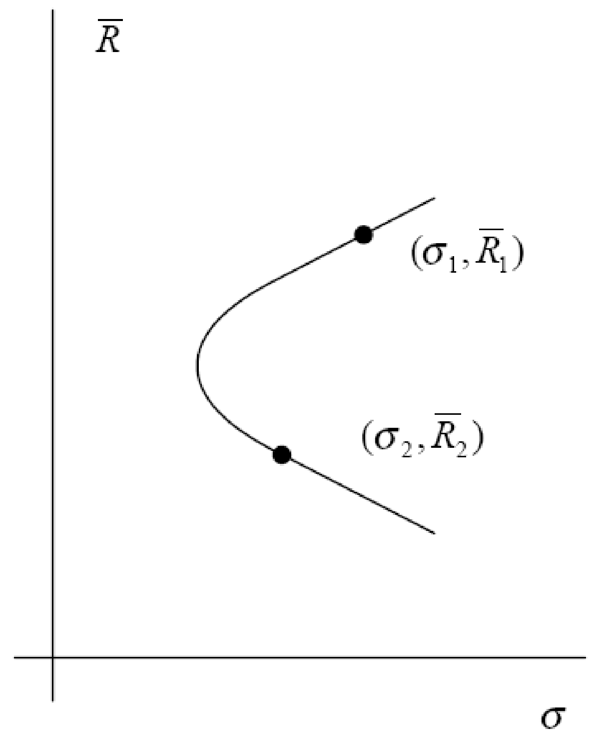

Global Minimum-Variance Portfolio

To find the *global minimum-variance portfolio* with the smallest possible variance, set the derivative to zero to get

$$w_{G1} = \frac{\sigma_2^2 - \sigma_{12}}{\sigma_1^2 + \sigma_2^2 - 2\sigma_{12}}, \qquad w_{G2} = \frac{\sigma_1^2 - \sigma_{12}}{\sigma_1^2 + \sigma_2^2 - 2\sigma_{12}}.$$

Portfolio Standard Deviation

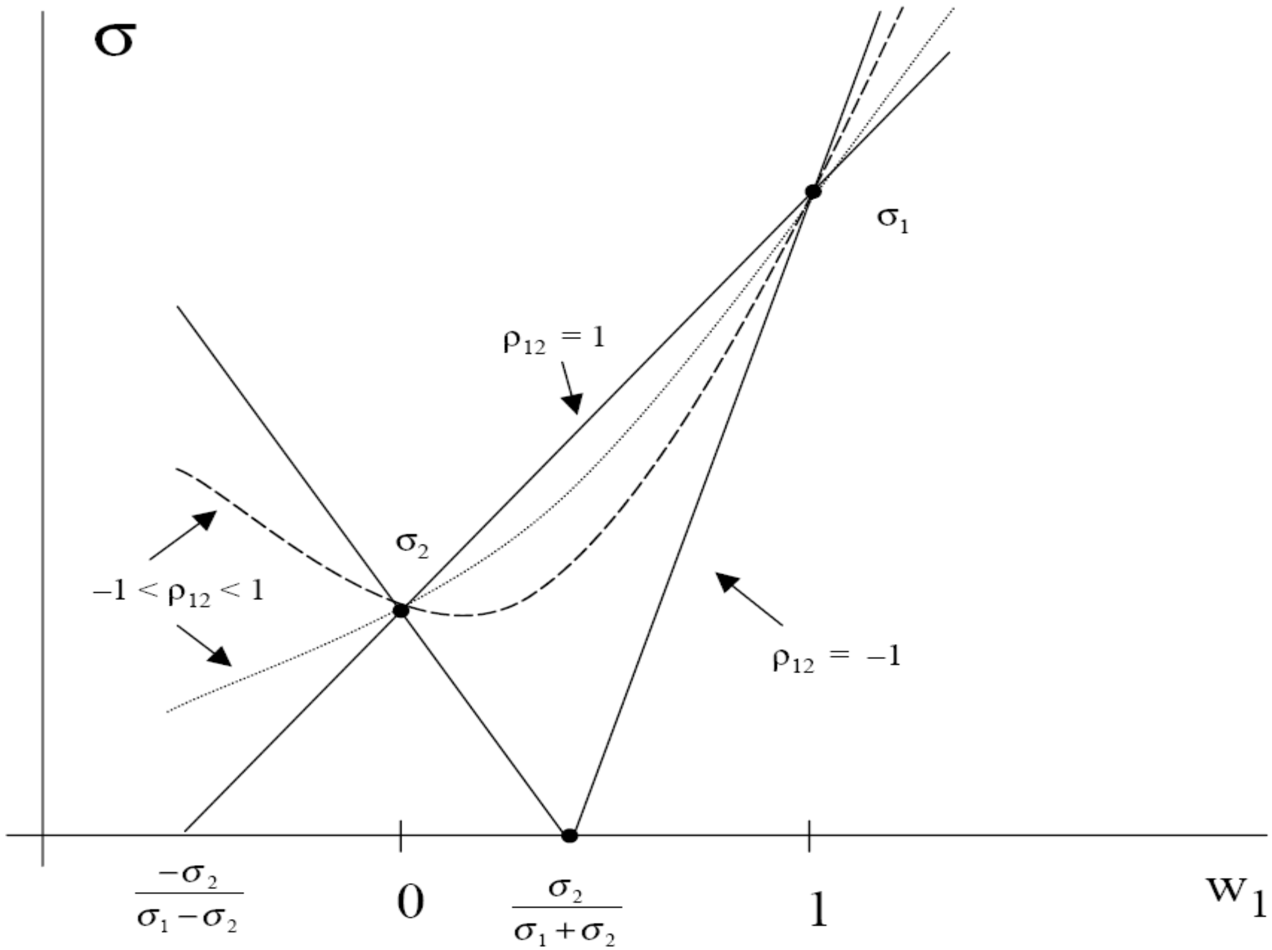

Global Minimum-Variance Portfolio

Special cases:

1. When $\sigma_{12} = 0$, $w_{G1} = \sigma_2^2/(\sigma_1^2 + \sigma_2^2)$, $\frac{\sigma_2^2}{\sigma_1^2+\sigma_2^2} w_{G2} = \sigma_1^2/(\sigma_1^2 + \sigma_2^2)$.
2. When $\sigma_1^2 = \sigma_2^2$, $w_{G1} = 1/2$, $w_{G2} = 1/2$.
3. When $\sigma_{12} = \sigma_1\sigma_2$ (so $\rho_{12} = 1$), the portfolio variance can be set to zero by

$$w_{G1} = \frac{-\sigma_2}{\sigma_1 - \sigma_2}, \qquad w_{G2} = \frac{\sigma_1}{\sigma_1 - \sigma_2}.$$

4. When $\sigma_{12} = -\sigma_1\sigma_2$ (so $\rho_{12} = -1$), the portfolio variance can be set to zero by

$$w_{G1} = \frac{\sigma_2}{\sigma_1 + \sigma_2}, \qquad w_{G2} = \frac{\sigma_1}{\sigma_1 + \sigma_2}.$$

One Risky, One Safe Asset

In this case $\sigma_2^2 = 0$ and $R_2 = R_f$, the riskless interest rate. We have $\overline{R}_p - R_f = w_1(\overline{R}_1 - R_f)$ and $\sigma_p^2 = w_1^2\sigma_1^2$ or $w_1 = \sigma_p/\sigma_1$. Hence

$$\overline{R}_p - R_f = \sigma_p \left(\frac{\overline{R}_1 - R_f}{\sigma_1} \right).$$

This defines a straight line, called the *capital allocation line (CAL)*, on a mean-standard deviation diagram.

Capital Allocation Line

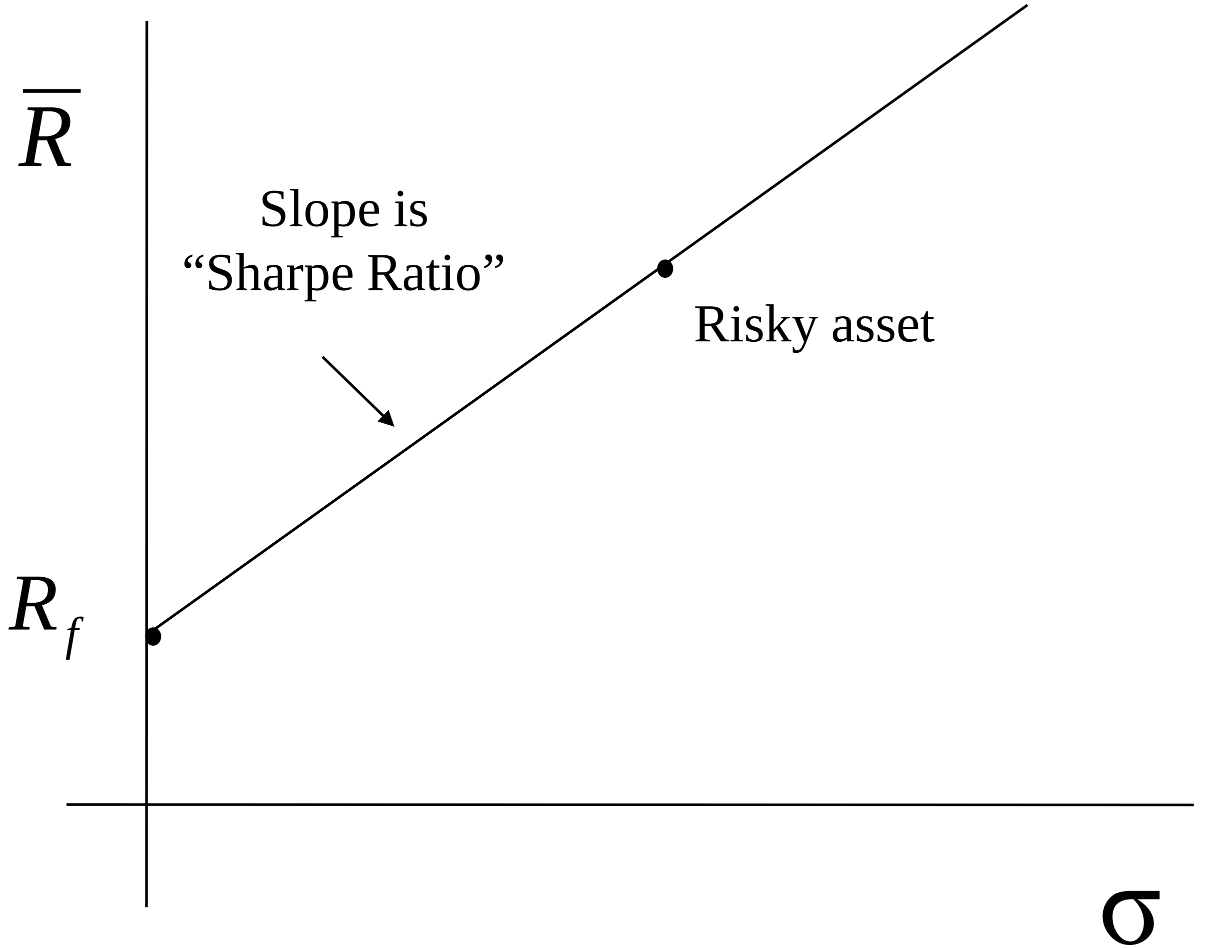

Sharpe Ratio

The slope

$$S_1 = \left(\frac{\overline{R}_1 - R_f}{\sigma_1} \right)$$

is called the *Sharpe ratio* of the risky asset. Any portfolio that combines a single risky asset with the riskless asset has the same Sharpe ratio as the risky asset itself.

The standard rule of myopic portfolio choice is

$$w_1 = \frac{\overline{R}_1 - R_f}{RRA\sigma_1^2} = \frac{S_1}{RRA\sigma_1}.$$