The Sharpe Ratio Ratio

Thomas Smith

June 18, 2019

$$\frac{\mathbf{w_2}}{\mathbf{w_1}} = \frac{\sigma_\mathbf{1}}{\sigma_\mathbf{2}} \frac{\mathbf{SRR} - \rho}{\mathbf{1} - \mathbf{SRR} \cdot \rho}$$

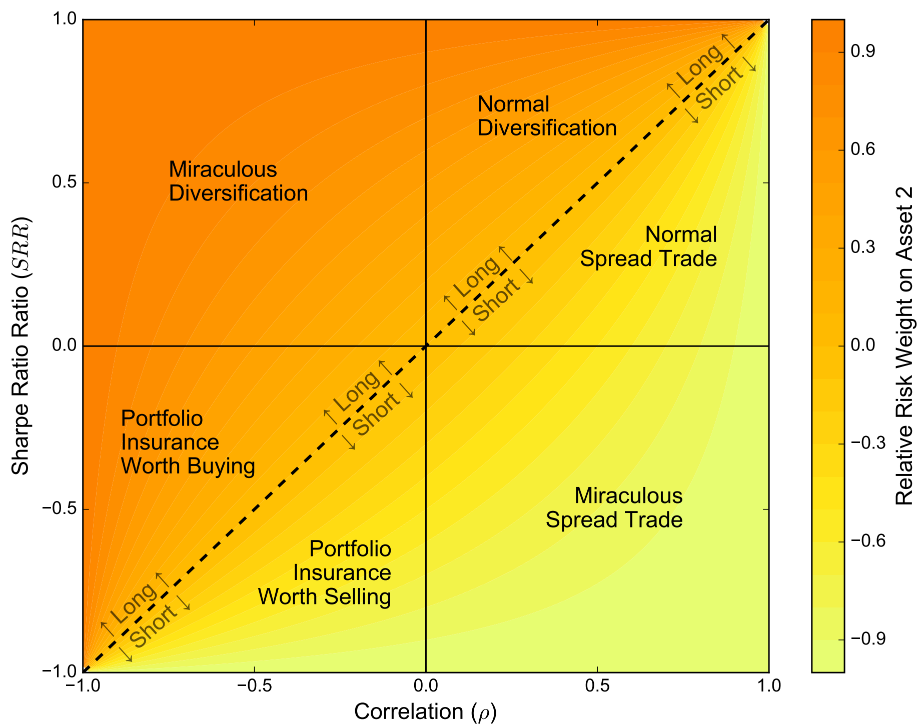

Abstract

Portfolio optimisers often give results that are extreme and unstable. With so many variables involved, it can be hard to decipher why. This article presents a simplified formula for a two-asset optimisation which will bolster your intuitive understanding of how a portfolio optimiser behaves and why it may behave badly.

We reduce the two-asset optimisation to a simple formula of two variables: the correlation and the ratio of the two assets' Sharpe Ratios. These variables are enough to illustrate why a portfolio optimiser can be so unstable.

While all the mathematical results are derived for a two-asset optimisation, the intuitive ideas can still be applied to n-assets.

Author

Thomas Smith specialises in quantitative investing. He currently works in the Quantitative Analytics division of Barclays plc. Previously, he was Lead Quant Researcher and partner at Brooksbridge Capital LLP, and a member of the fixed income research team at AHL, a division of Man Group plc. Thomas holds degrees in economics from the University of Toronto and the University of Oxford.

1 Extremist and Unstable

Portfolio optimisers are elegant, useful and - by definition - optimal. At least in theory. Anyone that has used one in practice can attest they are:

- **Extremist** The 'optimal' weights they output are often extremely unbalanced: massively long one asset, massively short another
- **Unstable** If you slightly change one input variable, the optimal weights shift dramatically

Why is it that with some inputs an optimiser behaves nicely and gives logical results, but with other inputs it seems to go crazy?

To better understand this behaviour, we will look at a simplified formula for a two-asset optimisation.

1.1 The Formula

These seemingly unpredictable behaviours can be understood by understanding Equation 1.

$$\frac{w_2}{w_1} = \frac{\sigma_1}{\sigma_2}\frac{SRR - \rho}{1 - SRR \cdot \rho} \tag{1}$$

This formula describes the optimal weights of a two-asset portfolio when we maximise the portfolio Sharpe Ratio. This tells us the optimal weights are driven by three inputs: the asset volatilities (σ_i), the correlation (ρ) and the Sharpe Ratio Ratio (SRR). This is the ratio of the Sharpe Ratios of the two assets ($SRR = \frac{SR_2}{SR_1}$).

On the left side of this equation, we have a fraction: $\frac{w_2}{w_1}$. This formula tells us the optimal *relative weight*. This relative weight is optimal because it maximizes the Sharpe Ratio of our portfolio.

1.2 Deriving the formula

To derive Equation 1 we start with the formula for the optimal portfolio weights, represented in matrices.

$$\mathbf{w_{tgt}} = \psi \cdot \mathbf{S^{-1}} \cdot \mathbf{C^{-1}} \cdot \mathbf{sr} \tag{2}$$

Equation 2 gives the optimal weights for a portfolio of n risky assets.

$\mathbf{w_{tgt}}$ and $\mathbf{sr}$ are vectors of the portfolio weights and asset Sharpe Ratios, respectively. $\mathbf{C}$ is the correlation matrix. And $\mathbf{S}$ is a matrix with the asset standard deviations along the diagonal and zeros for all off-diagonal elements.

$\mathbf{w}$ gets the subscript $\mathbf{tgt}$ because these weights achieve our target return.

ψ is a scalar that is included to scale the weights up or down to achieve the portfolio's target return.[1] You can find more detail about this scalar in Appendix 10.1. Because we are only interested in the relative weight of the portfolio, this scalar will cancel-out.

If we look at a portfolio of two-assets, then the inputs look like this:

$$\mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}, \mathbf{S} = \begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}, \mathbf{sr} = \begin{bmatrix} SR_1 \\ SR_2 \end{bmatrix}$$

Rewrite Equation 2 to solve for the optimal risk weights.

$$\mathbf{S} \cdot \mathbf{w} = \psi \cdot \mathbf{C}^{-1} \cdot \mathbf{sr}$$

We use the general formula for the inverse of a 2x2 matrix.

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad - bc} \cdot \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$$

Apply this operation to the correlation matrix, $\mathbf{C}$.

$$\begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{bmatrix} \cdot \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \psi \cdot \frac{1}{1 - \rho^2} \begin{bmatrix} 1 & -\rho \\ -\rho & 1 \end{bmatrix} \cdot \begin{bmatrix} SR_1 \\ SR_2 \end{bmatrix}$$

Multiply the matrices on either side.

$$\begin{bmatrix} w_1 \cdot \sigma_1 \\ w_2 \cdot \sigma_2 \end{bmatrix} = \frac{\psi}{1 - \rho^2} \cdot \begin{bmatrix} SR_1 - \rho \cdot SR_2 \\ -\rho \cdot SR_1 + SR_2 \end{bmatrix}$$

Because we are interested in the *relative* risk weight, we convert the two rows of our matrices into ratios.

$$\frac{w_2 \cdot \sigma_2}{w_1 \cdot \sigma_1} = \frac{\frac{\psi}{1-\rho^2} \cdot (SR_2 - \rho \cdot SR_1)}{\frac{\psi}{1-\rho^2} \cdot (SR_1 - \rho \cdot SR_2)}$$

[1] If the weights are scaled to achieve a target return of r_{tgt}, then the scalar ψ is equal to $\frac{r_{tgt} - r_f}{SR_{max}^2}$. If the weights are scaled to achieve a target volatility of σ_{tgt} then the scalar ψ is equal to $\frac{\sigma_{tgt}}{SR_{max}}$. For more details see Equations 9 and 10 in Appendix 10.1.

On the right side, divide both top and bottom by SR_1. Move the σ_i's over.

$$\frac{w_2}{w_1} = \frac{\sigma_1}{\sigma_2}\frac{SRR - \rho}{1 - SRR \cdot \rho}$$

This gives us Equation 1: our formula for the optimal relative weight of a two-asset portfolio. In the following sections we will use this formula to examine the optimiser's behaviour.

2 What do volatilities matter?

Our formula can be rearranged as follows:

$$\frac{w_2 \cdot \sigma_2}{w_1 \cdot \sigma_1} = \frac{SRR - \rho}{1 - SRR \cdot \rho} \tag{3}$$

The left side of Equation 3 gives us the *relative risk weight* of our optimal portfolio. The risk weight is our allocation to an asset, multiplied by the volatility of that asset.[2]

If the volatility of Asset 2 increased, but its Sharpe Ratio stayed the same (so SRR also stays the same), w_2 would decrease so that the risk weight of Asset 2 remains the same.

You can think of Equation 1 as a two-step process: the correlation and the Sharpe Ratio Ratio determine the optimal risk weights, and then the asset volatilities transform those risk weights into regular portfolio weights. The asset volatilities scale the weights up and down, ensuring we give less weight to high vol assets and more weight to low vol assets.

But it is the correlation and the Sharpe Ratio Ratio that provide all the interesting dynamics of mean-variance optimisation.

[2]Some investors refer to these as *vol-normalised weights* or *vol-scaled weights*.

3 I'm trying to build a long-only portfolio. Why does my optimiser give me erroneous spread trades?

Here is a problem, commonly encountered in practice: you take a pair of stocks, say *Ford Motor Co.* and *General Motors*, both have positive expected returns, reasonable volatilities and a positive correlation. You put them into an optimiser expecting it to return a sensible portfolio, something like 60%/40% between Ford/GM. Instead it gives you something like +180%/-80%: a huge long-position in Ford and a huge short-position in GM.

Ford and GM are in the same industry and face similar future prospects; a sensible portfolio should have similar risk allocations to each. Why does the optimiser want a spread-trade between the two?

The answer jumps out from Equation 3. A spread trade means our risk weights have opposite signs, so their ratio is negative.

$$\frac{w_2 \cdot \sigma_2}{w_1 \cdot \sigma_1} = \frac{SRR - \rho}{1 - SRR \cdot \rho} < 0$$

If we set our assets so that $|SR_2| \leq |SR_1|$ ('Asset 1' and 'Asset 2' are arbitrary monikers, we can always swap them to make this true), then we guarantee that $-1 \leq SRR \leq +1$. This also guarantees that $1 - SRR \cdot \rho \geq 0$, which gives:

$$SRR - \rho < 0$$

$$SRR < \rho$$

The optimiser wants a spread trade whenever the correlation is greater than the Sharpe Ratio Ratio.

A high correlation means two assets behave similarly; similar assets should have similar risk adjusted returns. If two assets behave similarly (high ρ) but with a large difference in returns (low SRR), then the optimiser wants to go long/short between them.

This also tells us when an optimiser will give us 'normal diversification' (meaning positive allocations to both assets, like we're usually trying to achieve). This will only occur if the correlation between the assets is less than their Sharpe Ratio Ratio ($\rho < SRR$).

4 I changed one of my expected returns from 6% to 5%. Why did the optimiser change my allocations so massively?

You tweaked your correlation estimate for Ford and GM so now instead of a spread trade your optimiser gives you two long positions. But its relative weight still has 98% in Ford and only 2% in GM. Way off balance. You decide to reduce your expected return for Ford by 1% per year. Now, the optimiser tells you to put only 7% in Ford and 93% in GM.

Why the wild change?

Mean-variance optimisation is often extremely sensitive to the input parameters. Equation 3 can tell us why. The denominator is $1 - SRR \cdot \rho$. If $SRR \cdot \rho \simeq 1$, then our denominator is $\simeq 0$.

With the denominator close to zero, the value of the numerator is magnified. Slight changes in ρ or SRR will lead to massive changes in our relative weight.

This is the danger of using an optimiser to combine similar assets. Similar means $\rho \simeq 1$ and $SRR \simeq 1$, which is precisely the case when our denominator $\simeq 0$. Even if the optimiser produces weights with the sign we want ($w_1 > 0, w_2 > 0$), with slight changes in our estimates of ρ or SRR our allocations will change wildly.

If the inputs for our assets change just enough so that ρ goes from being slightly less than SRR to slightly more than SRR, we will see our allocations swing from two long positions to a spread trade. All the result of a tiny change in our inputs.

5 I'm too lazy to deal with correlations. What should I do?

As we saw above, a high correlation makes your optimiser extremely sensitive to your inputs, specifically the SRR. Since you must estimate SR_2 and SR_1, they are likely to change. How can you make your portfolio more robust to these changes?

One simple way is to just ignore the correlation; pretend it is zero. What does this lead to?

$$\rho = 0 \Rightarrow \frac{w_2 \cdot \sigma_2}{w_1 \cdot \sigma_1} = \frac{SRR - \rho}{1 - SRR \cdot \rho} = \frac{SRR - 0}{1 - SRR \cdot 0} = \frac{SRR}{1} = SRR = \frac{SR_2}{SR_1}$$

If you assume $\rho = 0$ then an optimiser will give each asset a relative risk weight proportional to its Sharpe Ratio.

So if you have Sharpe Ratio estimates and you're happy to assume that correlation is zero, then you don't need an optimiser. You've already calculated the relative risk weight it would give you.

6 I'm also too lazy to deal with Sharpe Ratio estimates. What should I do?

If you're too lazy to estimate Sharpe Ratios, then you can just assume that the Sharpe Ratios are the same, which is equivalent to assuming $SRR = 1$.

$$SRR = 1 \Rightarrow \frac{w_2 \cdot \sigma_2}{w_1 \cdot \sigma_1} = \frac{SRR - \rho}{1 - SRR \cdot \rho} = \frac{1 - \rho}{1 - 1 \cdot \rho} = \frac{1 - \rho}{1 - \rho} = 1 \Rightarrow w_1 \cdot \sigma_1 = w_2 \cdot \sigma_2$$

If $SRR = 1$, then $w_2 \cdot \sigma_2 = w_1 \cdot \sigma_1$. You give equal risk weight to both assets, no matter what their correlation. As Ray Dalio will tell you, this is known as a *Risk Parity Portfolio*.[3]

7 Is diversifiying always better?

Looking at Equation 1, we can see one peculiar case: when the correlation exactly equals the Sharpe Ratio Ratio.

[3] Technically, this is an *Inverse Volatility Portfolio*. A Risk Parity Portfolio tries to equalise each asset's contribution to the portfolio's risk, which depends not only on the assets' volatilities but also their correlations. In the special case of two assets, where there is only one correlation involved, Inverse Volatility and Risk Parity happen to be the same thing.

$$\rho = SRR \Rightarrow \frac{w_2 \cdot \sigma_2}{w_1 \cdot \sigma_1} = \frac{SRR - \rho}{1 - SRR \cdot \rho} = \frac{0}{1 - SRR \cdot \rho} = 0$$

This surprised me. Just because $\rho = SRR$, we allocate nothing to Asset 2? What does this mean?

Let's assume we have Asset 1 with $SR_1 = 1$. We also have Asset 2 with $SR_2 = .5$ and its correlation with Asset 1 is $\rho = .5$. My first guess would be that Asset 2 is a reasonable diversifier to Asset 1. Asset 2 has a lower risk adjusted return, but given its low correlation I would expect it to get some allocation in an optimal portfolio. However, according to the optimiser, Asset 2 can't improve your portfolio.

When you add Asset 2, it has two effects: its low correlation provides a diversification benefit and reduces the portfolio volatility; but its low Sharpe Ratio reduces your portfolio's return. When $\rho = SRR$, these two effects offset each other exactly. No matter what your allocation to Asset 2, your portfolio Sharpe Ratio remains constant (and equal to SR_1).

8 Let's plot this

Let's imagine that our portfolio is already long Asset 1. Equation 3 gives us the optimal risk weight in Asset 2, expressed as a fraction of our risk weight in Asset 1 ($\frac{w_2 \cdot \sigma_2}{w_1 \cdot \sigma_1}$). This fraction is bound between -1 and +1 because we assume $|SR_2| \leq |SR_1|$.[4]

Figure 1 is a contoured heatmap where the intensity of the colour reflects the value of Equation 3. As we go from left to right the correlation goes from -1 to +1. And from bottom to top, the Sharpe Ratio Ratio goes from -1 to +1. And because Equation 3 is bound between -1 and +1, the heatmap encompasses all possible portfolios.

What does this plot show us?

8.1 The 45° line

The dashed 45° line is where $\rho = SRR$, which is the dividing line between going long/short Asset 2. To the right of the 45° line ($\rho > SRR$), the

[4] As I said before, 'Asset 1' and 'Asset 2' are arbitrary monikers, so we can always swap them around to make this assumption true.

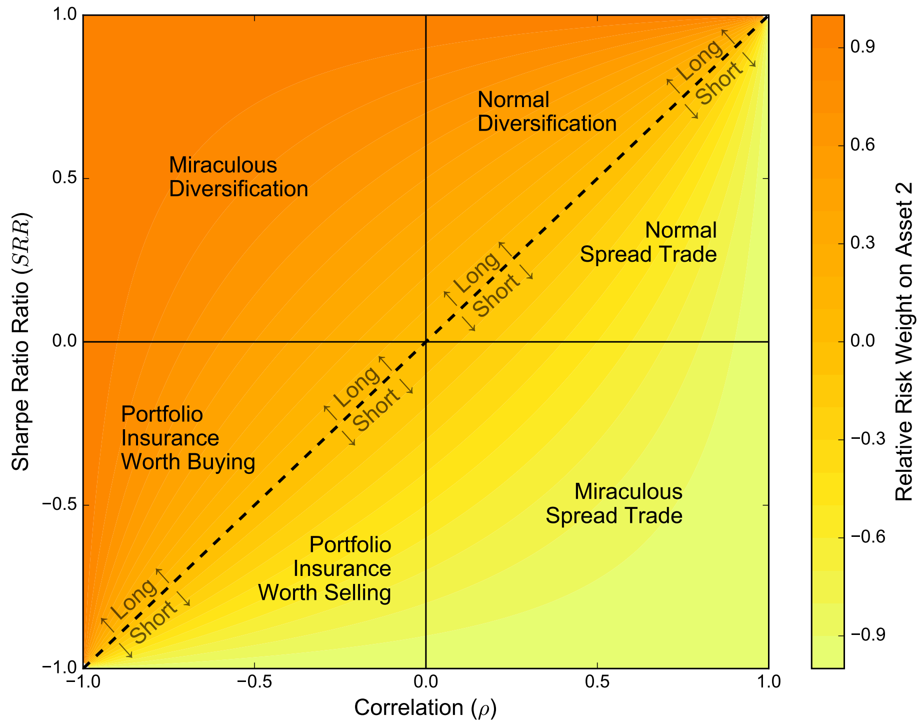

Figure 1: Heatmap of the Relative Risk Weight on Asset 2 as a Function of the Correlation and the Sharpe Ratio Ratio

optimiser wants a spread trade. To the left of the 45° line ($\rho < SRR$), the optimiser wants to be long both assets. Close to the 45° line, the risk weight that the optimiser gives to Asset 2 is close to zero.

Far above the 45° line, where the heatmap is darkest, the optimiser gives a large and positive risk weight to Asset 2. Along the top line $SRR = 1$ and the optimiser goes long Asset 2 with a risk weight equal to Asset 1.

Far below the 45° line, where the heatmap is lightest, the optimiser gives a large and negative risk weight to Asset 2. Along the bottom line $SRR = -1$ and the optimiser goes short Asset 2 with a risk weight equal in absolute magnitude to Asset 1.

8.2 The most important quadrant

The region of greatest interest is the top right quadrant, where $SRR > 0$ and $\rho > 0$. In practice, you will mostly be dealing with assets that have positive correlations and positive returns (hence positive Sharpe Ratios and positive Sharpe Ratio Ratios).

Looking in this region, the contours of the heatmap illustrate the most frustrating behaviour of mean-variance optimisation.

Along the vertical middle line, where the correlation is zero, we can see that all the contours are equally thick. This reflects that when $\rho = 0$ the optimal risk weight on Asset 2 will be proportional to SRR.

As we move from left to right in this quadrant, all the contours bunch together. This is because as the correlation increases, the optimiser becomes more and more sensitive to the Sharpe Ratio Ratio.

Near the right edge of this quadrant, where $\rho \simeq 1$, all the contours converge. In this area, if we move up or down slightly, we can cross almost every contour on the map. This reflects that when ρ is high, the optimiser is extremely sensitive to any slight change in SRR.

The converse is not true. Looking along the top line, we can see that when the SRR is close to $+1$, the optimiser is not sensitive to changes in ρ. No matter what the correlation, the optimiser gives Asset 2 a relative risk weight $\simeq 1$.

8.3 Naming the regions

On the plot, we've applied names to each region which relate ρ and SRR to different investing scenarios.

- **Normal Diversification** ($0 < \rho < SRR$)

  Long-only combinations of positively correlated assets, where both assets have positive expected returns (hence positive Sharpe Ratios and positive Sharpe Ratio Ratios).

  This is *normal* diversification because, in practice, most investors trying to build portfolios will be working in this region.

- **Normal Spread Trade** ($0 < SRR < \rho$)

  If you work at a hedge fund you might be doing more complicated trading strategies where you want to put on a spread trade: going long/short in Asset 1/Asset 2.

  In a normal spread trade, both assets have positive expected returns but you still expect to gain because, on a risk adjusted basis, you expect Asset 2 to return less than Asset 1 ($SRR < 1$). And because the assets have a positive correlation, your long/short positions will reduce your portfolio volatility.

  As we discussed, a mean-variance optimiser often finds erroneous spread trades: you are trying to build a portfolio with normal diversification but because you have estimates where $\rho > SRR$ the optimiser gives you a spread trade that you did not intend.

- **Miraculous Diversification** ($\rho < 0 < SRR$)

  If you're a long-only investor, your dream is to find an asset that has a positive Sharpe Ratio and a negative correlation to your existing portfolio. This is *miraculous* diversification because it usually doesn't exist.

  If all you hold is stocks, then you can probably get some miraculous diversification by adding bonds, but even the correlation of stocks and bonds is not deeply negative; it is usually close to zero and has varied over time between positive and negative. If you can find an asset with positive returns and negative correlation, it is a great diversifier and you should definitely buy it.

  Most of the area in this region is coloured very dark, which means that for most values of ρ and SRR you would give Asset 2 a risk weight nearly equal to Asset 1. This also means that in this region the optimiser is not very sensitive, rather it is robust. You could change the values of ρ and SRR and your optimal weights would not change much.

- **Miraculous Spread Trade** ($SRR < 0 < \rho$)

  This is like miraculous diversification because it usually doesn't exist, but if you can find it you should definitely take it.

This region represents assets that have a positive correlation, but Asset 1 has a positive expected return while Asset 2 has a negative expected return. If you go long Asset 1 and short Asset 2, you expect to gain on both sides of the trade. And because of their positive correlation, the long/short position will reduce your portfolio vol, giving your portfolio a great risk adjusted return.

- **Portfolio Insurance Worth Buying** ($\rho < SRR < 0$)

  You may find that Asset 2 has a negative expected return, but its correlation with Asset 1, which you already hold, is so strongly negative that it is still worth buying Asset 2.

  Assets like these are referred to as 'portfolio insurance'; think of a tail-risk protection fund. On average, you expect to lose money on them, but their negative correlation reduces your portfolio vol so much that they improve your portfolio's risk adjusted return (your portfolio Sharpe Ratio).

  Imagine you hold the S&P500, and the market price plummets for out-of-the-money S&P500 put options. You expect the S&P500 to go up, so the puts have a negative expected return. But the puts are a good hedge that reduce your portfolio vol. If you can buy them cheaply enough they will increase your portfolio Sharpe Ratio.[5]

- **Portfolio Insurance Worth Selling** ($SRR < \rho < 0$)

  You may find that portfolio insurance is available, but it is so expensive that it isn't worth buying (so you prefer to short it).

  Imagine the market price skyrockets for out-of-the-money S&P500 puts. The puts have a strong negative correlation with your portfolio, but they are so expensive, and their negative expected return is so great, you prefer to short them and earn the option premium. You are selling portfolio insurance.

[5] This is a theoretical scenario. In practice, out-of-the-money put options are usually overpriced, so buying them would reduce your portfolio Sharpe Ratio.

9 What if I have more than two assets?

Focusing on the two-asset portfolio keeps our illustration neat and simple. But in the real world, messy and complex, you are almost always dealing with more than two-assets. What insights can we translate from the two-asset case to the n-asset case?

9.1 Just pretend

While only a foolish investor holds a two-asset portfolio, a clever investor may still pretend they do.

If you are considering making a new investment, but without changing any of your existing investments, you can pretend you are dealing with just two assets: Asset 1 is your *Existing Portfolio*; Asset 2 is your *New Investment*. Then all the two-asset insights we have covered still apply.

There are scenarios where this simplification is sufficiently realistic: if all you hold is stocks and you are considering adding bonds to your portfolio. Or if you are considering a portfolio-insurance investment; you would not be concerned with the correlation of the portfolio-insurance with any individual asset, only its correlation with your portfolio as a whole.

9.2 Replace 'correlation' with 'asset redundancy'

A two-asset optimiser is extremist and unstable when the correlation is high. An n-asset optimiser is extremist and unstable, but rather than depending on a single correlation, it depends on the asset's redundancy.

'Asset Redundancy' is the extent to which one asset's returns are a linear combination of other assets. Think of a portfolio that contains a stock ETF, a bond ETF and a balanced mutual fund which holds a 60/40 portfolio of stocks and bonds. The balanced mutual fund is largely redundant, as it is just a combination of the other assets.[6]

If you have one asset that is entirely redundant (a perfect linear combination of other assets) then your correlation matrix will be singular, thus non-invertible, thus the whole optimisation will throw an error.

[6] In an n-asset portfolio, the simplest case of asset redundancy is just two assets with a high correlation. If one is included, the other is largely redundant.

We can actually take a numeric measure of an asset's redundancy: do a linear regression of the returns of $Asset_i$ on the returns of all other assets in your optimisation; the R^2 from that regression is a measure of how redundant $Asset_i$ is.

If the R^2 is near-zero, then $Asset_i$ is mostly uncorrelated with your other assets and thus a good diversifier. It is not redundant at all. If the R^2 is exactly 1.0, then $Asset_i$ is a perfect linear combination of the other assets, and thus optimisation is impossible.

The frustrating situation is dealing with an asset that is highly redundant, but not perfectly so: the R^2 is high, but less than 1.0 (something like $0.8 \leq R^2 < 1.0$). This is when your optimiser is is likely to be extremist and unstable.

9.2.1 The intuition is analogous

The intuition we covered for a two-asset optimisation with highly correlated assets, still applies to an n-asset optimisation. But instead of highly correlated assets, it applies when we have a highly redundant asset.

In an n-asset optimisation, when $Asset_i$ is highly redundant, the optimisation is highly sensitive to SR_i. A slight change in SR_i will lead to a huge change in w_i.[7]

And you can still get an erroneous short position in $Asset_i$ if SR_i is sufficiently low. How low? That is hard to define precisely: as well as depending on SR_i, it depends on the Sharpe Ratios of each of the other assets, and the correlation of $Asset_i$ with each of the other assets.

In the two-asset case, we could show precisely when w_2 will be negative: when $SRR < \rho$. In the n-asset case, the behaviour is more complex but the intuition is analogous: if $Asset_i$ is highly redundant, and SR_i is sufficiently low relative to the Sharpe Ratios of the other assets, then w_i will be negative.

Unfortunately, giving a precise definition of the phrase 'sufficiently low' is beyond the mathematical detail we will go into here.

[7] The sensitivity of an optimisation to slight changes in the Sharpe Ratio estimates can be described mathematically by the condition number of the correlation matrix. A set of assets with high multicollinearity will have a high condition number.

10 Appendix

10.1 Deriving the optimal weights

In this section we derive Equation 2: our equation for the optimal weights of an n-asset portfolio. This derivation largely follows Zakamulin, 2011 [2].

10.1.1 The setup

r_f is the risk-free rate. r_i is the return on asset i. e_i is the excess return of asset i over the risk-free rate, so $e_i = r_i - r_f$. w_i is the portfolio weight allocated to asset i.

$\mathbf{w}$ is the vector of weights. $\mathbf{r}$ is the vector of asset returns. $\mathbf{e}$ is the vector of excess returns. $\mathbf{\Sigma}$ is the covariance matrix.

$$\mathbf{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}, \mathbf{r} = \begin{bmatrix} r_1 \\ \vdots \\ r_n \end{bmatrix}, \mathbf{e} = \mathbf{r} - r_f = \begin{bmatrix} r_1 - r_f \\ \vdots \\ r_n - r_f \end{bmatrix}, \mathbf{\Sigma} = \begin{bmatrix} \sigma_1^2 & \dots & \sigma_{1,n} \\ \vdots & \ddots & \vdots \\ \sigma_{n,1} & \dots & \sigma_n^2 \end{bmatrix}$$

We assume that we can borrow and lend all we want at the risk-free rate, r_f. w_f is the portfolio weight allocated to the risk-free asset; it represents how much we borrow/lend.

$$w_f = 1 - \sum_{i=1}^{n} w_i$$

The weight given to the risk-free asset, plus the weights of the n risky assets, will always sum to one.

The return on our portfolio is our borrowing/lending times the risk-free rate, plus our allocations to the risky assets times their respective returns.

$$r_p = w_f \cdot r_f + \sum_{i=1}^{n} w_i \cdot r_i$$

Substitute in our equation for w_f and we can express our portfolio return in terms of our vector of excess returns.

$$r_p = (1-\sum_{i=1}^{n} w_i)\cdot r_f+\sum_{i=1}^{n} w_i\cdot r_i = r_f+\sum_{i=1}^{n} w_i\cdot(r_i-r_f) = r_f+\sum_{i=1}^{n} w_i\cdot e_i = r_f+\mathbf{w}^T\cdot\mathbf{e}$$

10.1.2 The optimisation

For our optimisation, we start with the formula for the volatility of our portfolio, σ_p^2.

$$\sigma_p^2 = \mathbf{w}^T\cdot\mathbf{\Sigma}\cdot\mathbf{w}$$

We will minimise our portfolio volatility, subject to the constraint that our portfolio return equals our target return, r_{tgt}.[8]

$$\min_{\mathbf{w}} \frac{1}{2}\mathbf{w}^T\cdot\mathbf{\Sigma}\cdot\mathbf{w} \quad s.t. \quad r_p = r_f + \mathbf{w}^T\cdot\mathbf{e} = r_{tgt} \tag{4}$$

Figure 2 is a stylized depiction of our portfolio efficient frontiers. Because our optimisation includes the risk-free asset, we are solving directly for the straight line: the Capital Market Line.

This means that, while we set up our optimisation to minimise our portfolio volatility, our result gives us a portfolio with the maximum achievable Sharpe Ratio (SR_{max}). This is the highest Sharpe Ratio that it is possible to achieve with this set of assets.

This setup reflects the Two-Fund Separation Theorem, which says that investors should hold a combination of the risk-free asset and the portfolio of risky assets with the maximum achievable Sharpe Ratio. The Sharpe Ratio is invariant under leverage, so it stays constant no matter what our target return.

We solve this optimisation using a Lagrangian.

$$L = \frac{1}{2}\mathbf{w}^T\cdot\mathbf{\Sigma}\cdot\mathbf{w} + \lambda\cdot(r_f + \mathbf{w}^T\cdot\mathbf{e} - r_{tgt})$$

First, we take the derivative with respect to $\mathbf{w}$.

[8] In Equation 4, you can see that we actually minimise $\frac{1}{2}\sigma_p^2$. The $\frac{1}{2}$ scalar is included just for convenience as it will cancel-out later and minimising one half of the portfolio volatility will give the same result as minimising the entire portfolio volatility.

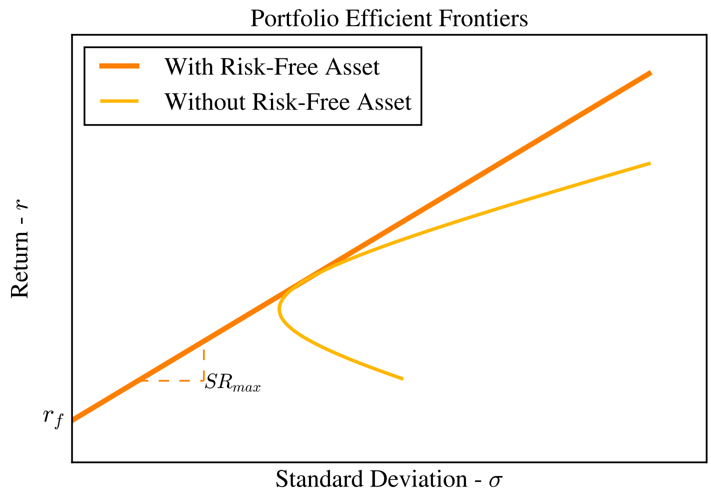

Figure 2: Efficient Frontiers

$$\frac{\delta L}{\delta w} = \frac{1}{2} \cdot \mathbf{\Sigma} \cdot \mathbf{w} \cdot 2 + \lambda \cdot \mathbf{e}$$

We set this derivative to zero. $\mathbf{w_{tgt}}$ is the label we give to the optimal weights that achieve our target return.

$$\frac{\delta L}{\delta w} = \frac{1}{2} \cdot \mathbf{\Sigma} \cdot \mathbf{w_{tgt}} \cdot 2 + \lambda \cdot \mathbf{e} = 0$$

Rearrange to solve for $\mathbf{w_{tgt}}$.

$$\mathbf{w_{tgt}} = -\lambda \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}$$

Second, we take the derivative with respect to the Lagrange multiplier and set this to zero.

$$\frac{\delta L}{\delta \lambda} = r_f + \mathbf{w_{tgt}}^T \cdot \mathbf{e} - r_{tgt} = 0$$

We substitute in our formula for $\mathbf{w_{tgt}}^T$.

$$\frac{\delta L}{\delta \lambda} = r_f + (-\lambda \cdot \mathbf{e}^T \cdot \mathbf{\Sigma^{-1}}) \cdot \mathbf{e} - r_{tgt} = 0$$

Rearrange to solve for $-\lambda$.

$$-\lambda = \frac{r_{tgt} - r_f}{\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}}$$

And substitute $-\lambda$ into our equation for $\mathbf{w_{tgt}}$.

$$\mathbf{w_{tgt}} = \frac{r_{tgt} - r_f}{\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}} \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e} \tag{5}$$

Equation 5 gives us the optimal weights that achieve our target return.

We could stop here and continue all our work with this formula. But the results in 10.1.3 will make this formula easier to interpret.

10.1.3 The highest achievable Sharpe Ratio

In this section, we will derive a formula for the Sharpe Ratio of our portfolio.

We start with the equation for the volatility of our portfolio when we use our optimal weights.

$$\sigma_{tgt}^2 = \mathbf{w_{tgt}}^T \cdot \mathbf{\Sigma} \cdot \mathbf{w_{tgt}}$$

We then substitute in Equation 5 for $\mathbf{w_{tgt}}$.

$$\sigma_{tgt}^2 = \left(\frac{r_{tgt} - r_f}{\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}} \cdot \mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \right) \cdot \mathbf{\Sigma} \cdot \left(\frac{r_{tgt} - r_f}{\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}} \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e} \right)$$

The fractions reduce to scalars, so we move both to the front. And one $\mathbf{\Sigma^{-1}}$ cancels with $\mathbf{\Sigma}$.

$$\sigma_{tgt}^2 = \frac{(r_{tgt} - r_f)^2}{(\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e})^2} \cdot \mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}$$

Cancel one of the denominators.

$$\sigma_{tgt}^2 = \frac{(r_{tgt} - r_f)^2}{\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}}$$

Rearrange.

$$\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e} = \left(\frac{r_{tgt} - r_f}{\sigma_{tgt}} \right)^2$$

$\frac{(r_{tgt}-r_f)}{\sigma_{tgt}}$ is the definition of the Sharpe Ratio. The left side of this equation is equal to the square of the Sharpe Ratio, and the square-root is equal to the Sharpe Ratio.

$$\sqrt{\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}} = \frac{r_{tgt} - r_f}{\sigma_{tgt}} = SR_{max} \tag{6}$$

This is the Sharpe Ratio of our portfolio when we use our optimal weights. Because our optimisation includes the risk-free asset, so we are solving for the straight line in Figure 2, this is also the highest possible Sharpe Ratio that can be achieved with these assets; we label this SR_{max}.

10.1.4 Target return

Let's go back to Equation 5; our formula for the optimal weights.

$$\mathbf{w_{tgt}} = \frac{r_{tgt} - r_f}{\mathbf{e}^T \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}} \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}$$

Substitute in Equation 6, our formula for the portfolio Sharpe Ratio.

$$\mathbf{w_{tgt}} = \frac{r_{tgt} - r_f}{SR_{max}^2} \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e} \tag{7}$$

This gives us the formula for our optimal portfolio weights, as a function of our target return.

10.1.5 Target volatility

Some investors use a target volatility, rather than a target return. We can restate Equation 7 in terms of a target volatility.

From the definition of the Sharpe Ratio we have

$$\frac{1}{SR_{max}^2} = \frac{\sigma_{tgt}^2}{(r_{tgt} - r_f)^2}$$

We substitute this into Equation 7.

$$\mathbf{w_{tgt}} = \frac{\sigma_{tgt}^2}{(r_{tgt} - r_f)^2} \cdot (r_{tgt} - r_f) \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}$$

Cancel one of the denominators.

$$\mathbf{w_{tgt}} = \frac{\sigma_{tgt}^2}{(r_{tgt} - r_f)} \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e}$$

Substitute in $1/SR_{max}$.

$$\mathbf{w_{tgt}} = \frac{\sigma_{tgt}}{SR_{max}} \cdot \mathbf{\Sigma^{-1}} \cdot \mathbf{e} \tag{8}$$

Equation 8 gives us our optimal portfolio weights as a function of a target volatility.

10.1.6 With correlations

In this section we will restate our optimal portfolio weights as a function of the asset Sharpe Ratios and correlation matrix.

We define three more variables. **C** is the correlation matrix. **S** is a diagonal matrix with the asset standard deviations across the diagonal and zeros for all other values. **sr** is the vector of Sharpe Ratios for each asset.

$$\mathbf{C} = \begin{bmatrix} 1 & \rho_{1,2} & \dots & \rho_{1,n} \\ \rho_{2,1} & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ \rho_{n,1} & \dots & \dots & 1 \end{bmatrix}, \mathbf{S} = \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & & \vdots \\ \vdots & & \ddots & \vdots \\ 0 & \dots & \dots & \sigma_n \end{bmatrix}$$

$$\mathbf{sr} = \begin{bmatrix} SR_1 \\ \vdots \\ SR_n \end{bmatrix} = \begin{bmatrix} \frac{r_1 - r_f}{\sigma_1} \\ \vdots \\ \frac{r_n - r_f}{\sigma_n} \end{bmatrix} = \begin{bmatrix} \frac{e_1}{\sigma_1} \\ \vdots \\ \frac{e_n}{\sigma_n} \end{bmatrix} = \mathbf{S^{-1}} \cdot \mathbf{e}$$

The covariance matrix, $\mathbf{\Sigma}$, is equivalent to

$$\mathbf{\Sigma} = \mathbf{S} \cdot \mathbf{C} \cdot \mathbf{S}$$

S and **C** are both symmetric, so the inverse of the covariance matrix is

$$\mathbf{\Sigma^{-1}} = \mathbf{S^{-1}} \cdot \mathbf{C^{-1}} \cdot \mathbf{S^{-1}}$$

Substitute this into Equation 7.

$$\mathbf{w_{tgt}} = \frac{r_{tgt} - r_f}{SR_{max}^2} \cdot \mathbf{S^{-1}} \cdot \mathbf{C^{-1}} \cdot \mathbf{S^{-1}} \cdot \mathbf{e}$$

Replace $\mathbf{S^{-1}} \cdot \mathbf{e}$ with $\mathbf{sr}$.

$$\mathbf{w_{tgt}} = \frac{r_{tgt} - r_f}{SR^2_{max}} \cdot \mathbf{S^{-1}} \cdot \mathbf{C^{-1}} \cdot \mathbf{sr} \tag{9}$$

This gives us the final formula for our optimal portfolio weights, in terms of a target return.

Applying the same steps to Equation 8 we get

$$\mathbf{w_{tgt}} = \frac{\sigma_{tgt}}{SR_{max}} \cdot \mathbf{S^{-1}} \cdot \mathbf{C^{-1}} \cdot \mathbf{sr} \tag{10}$$

This gives us the final formula for our optimal portfolio weights, in terms of a target volatility.

References

[1] Manuel Tarrazo. 'Portfolio Optimization with Correlation Matrices: How, Why, and Why Not'. *Journal of Finance and Investment Analysis*, 2(3), 2013. https://EconPapers.repec.org/RePEc:spt:fininv:v:2:y:2013:i:3:f:2_3_2.

[2] Valeriy Zakamulin. 'Sharpe (Ratio) Thinking About the Investment Opportunity Set and CAPM Relationship'. *Economics Research International*, Vol. 2011. https://ssrn.com/abstract=1479495.