The Geometry of Risk

Abstract

A geometric account of Markowitz's theory aimed at a student audience.

1 Pythagoras and Risk

What can geometry tell us about finance? A hint that geometry may have something interesting to say comes from the relationship between risk and Pythagoras's theorem.

To see this relationship, first we need to decide what we actually mean by risk. Quantifying the risk of an investment is a complex and controversial subject. The first serious attempt at quantifying risk mathematically was made by Markowitz in 1952 [1]. He was rewarded for his efforts with the 1990 Nobel Prize in Economics.

Let's follow Markowitz's approach, and quantify the risk of an investment by using the standard deviation of its payoff at a future time T.

If A and B are uncorrelated random variables with finite standard deviation, then their variances are related by the formula

$$\mathrm{Var}(A) + \mathrm{Var}(B) = \mathrm{Var}(A + B).$$

Therefore their standard deviations are related by the formula

$$\sigma(A)^2 + \sigma(B)^2 = \sigma(A + B)^2. \tag{1}$$

This looks suspiciously like Pythagoras's theorem for the length, $|\mathbf{a} + \mathbf{b}|$ of the sum of two perpendicular vectors $\mathbf{a}$ and $\mathbf{b}$:

$$|\mathbf{a}|^2 + |\mathbf{b}|^2 = |\mathbf{a} + \mathbf{b}|^{\mathbf{2}}.$$

The similarities between the properties of distance and the properties of standard deviation do not stop here: if we assume all non-trivial risk-free investments in the market have some risk, we can show that distance in Euclidean geometry and standard deviation in markets are mathematically identical. We will then be able to use geometry to obtain investment advice.

2 Diversification

Before tackling the full theory, let's examine an important special case which we will call *a market of identical stocks.* In a market of identical stocks, n stocks are traded and they have the following properties:

1. The units used to measure quantities of stocks are chosen so that initially each unit of stock costs \$1.

2. The stocks' payoffs are independent.

3. The stocks are financially equivalent to each one another in the sense that they all have the same probability distribution for their payoff at time T. This probability distribution has mean μ and standard deviation $\sigma > 0$.

If you have an amount, C, to invest in a market of identical stocks, how should you invest? You might guess that it wouldn't matter how you invested since the stocks are identical, but if you were to guess that, you would be wrong.

The trick is to realise it is possible to invest in a *portfolio* of stocks. You can invest an amount q_1 in stock 1, q_2 in stock 2, and so on. We represent this portfolio of stocks as a vector $\mathbf{q} = (q_1, q_2, \ldots q_n) \in \mathbb{R}^n$. We will see that investing in a portfolio of stocks rather than in just one specific stock allows you to reduce risk without reducing the mean payoff.

Let us calculate the mean payoff and its standard deviation. Write X_i for the random variable representing the payoff of stock i. The value of our portfolio at time T is a random random variable, Q, given by

$$Q = q_1 X_1 + q_2 X_2 + \ldots + q_n X_n. \tag{2}$$

Since one unit of stock has initial cost 1 and mean payoff μ, the cost, C, and mean payoff, P, of the portfolio are

$$\begin{aligned} C &= q_1 + q_2 + \ldots + q_n, \\ P &= (q_1 + q_2 + \ldots + q_n)\,\mu = C\,\mu. \end{aligned}$$

We conclude that the mean payoff for any portfolio is determined by C.

The standard deviation is more interesting. First note that the standard deviation of $q_i X_i$ is equal to $q_i \sigma$. We find from (2) and repeated applications of our analogue of Pythagoras's Theorem, equation (1), that the variance of the portfolio is given by

$$q_1^2\sigma^2 + q_2^2\sigma^2 + \ldots q_n^2\sigma^2 = |\mathbf{q}|^{\mathbf{2}}\sigma^{\mathbf{2}}.$$

We have shown:

Lemma. *The standard deviation of payoff of the portfolio corresponding to* $\mathbf{q}$ *in a market of identical stocks is proportional to the length of* $\mathbf{q}$. *The mean payoff is determined by the cost.*

Now consider a two dimensional market as illustrated in Figure 1. Each point in Figure 1 represents a portfolio with quantities (q_1, q_2). We have drawn lines indicating the cost of each portfolio. These lines are at 45 degrees because each stock in the market has the same cost: so the cost of the portfolio $(1, 0)$ is equal to the cost of the portfolio $(0, 1)$.

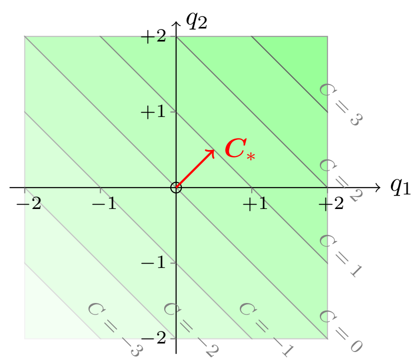

Figure 1: Contour lines of constant cost, C, when each stock costs \$1. The optimal portfolio of cost 1 is labelled $\mathbf{C}_*$

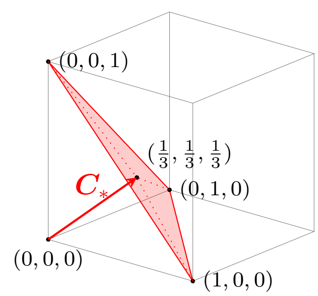

Figure 2: The optimum portfolio, $\mathbf{C}_*$ in a three dimensional market of identical assets

Among all the portfolios of cost \$1, which has the lowest risk? It is the portfolio on the line $C = 1$ closest to the origin, namely $\mathbf{C}_* = (\frac{1}{2}, \frac{1}{2})$, as shown in Figure 1. By Pythagoras's theorem, it has standard deviation

$$\left(\sqrt{\frac{1}{2}^2 + \frac{1}{2}^2}\right)\sigma = \frac{1}{\sqrt{2}}\sigma$$

We conclude that investing equal amounts in each stock is less risky than investing in just one stock, even though both portfolios have identical mean payoff.

In a general when investing \$1 in an n-dimensional market of identical stocks, you should divide your investment equally between all n stocks to give a portfolio

$$\mathbf{C}_* = \left(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\right). \tag{3}$$

This will have payoff with standard deviation $\frac{1}{\sqrt{n}}\sigma$. Dividing your investment across multiple assets in this way is called *diversification*.

The three dimensional case is illustrated in Figure 2. To see why this will be optimal, notice that the vector $\mathbf{C}_*$ given by (3) connects the origin perpendicularly to the hyperplane $C = 1$. So if $\mathbf{P}$ is another point on the hyperplane, we have by Pythagoras's theorem that:

$$|\mathbf{P}|^\mathbf{2} = |\mathbf{C}_*|^\mathbf{2} + |\mathbf{P} - \mathbf{C}_*|^\mathbf{2} \geq |\mathbf{C}_*|^\mathbf{2}.$$

We summarize our results.

Theorem. *When investing an amount C in a market of identical stocks, if you wish to minimize the risk of your investments (as measured by standard deviation) you should* diversify *your investment equally among all stocks. The mean payoff is determined entirely by the cost.*

This gives a mathematical explanation of the standard investment advice "don't put all your eggs in one basket". Giving a mathematical formulation of this simple, but important, result is a significant achievement of Markowitz's theory.

3 Financial markets and Euclid's axioms

In the general case, we will assume that we are given an n-dimensional set of random variables, V, which we call the *space of assets* or the *market.* Each random variable in V represents the payoff distribution of an asset or a portfolio of assets. We will assume that V contains all the assets in our market and all possible portfolios.

We will allow portfolios to include either positive or negative quantities of an asset. Holding negative quantities of an asset can be interpreted as promising to give someone that asset in the future. This generalizes the familiar idea of using negative quantities in your accounts to represent owing someone money. In finance taking a negative position in a stock is called *going short* on the stock. A positive position is called a *long* position. Our assumption that V contains all portfolios, including both long and short positions, can be expressed mathematically by saying that we assume V is a vector space.

We would like to interpret the standard deviation geometrically as the distance from the origin in this more general set up. To do this we will need to assume that all non-zero assets have a finite, non-zero standard deviation. If we make this assumption, it turns out that the standard deviation has all the same properties as distance in Euclidean geometry.

One way to understand why this might be true, is to recall that Euclid developed his theory of geometry by assuming only a very few basic axioms about points, lines and distances. If we can show that the distance in our market shares these same basic properties with Euclidean distance, it will follow that every theorem in Euclidean geometry will be equally true of the standard deviation in our market.

At this point in our argument you could go off and brush up your ancient Greek to find out what Euclid's axioms actually were. It is easier, however, to see Euclid's axioms as motivation and use modern linear algebra instead. From this modern point of view, Euclid's geometry is the theory of real inner product spaces. The assumption that the standard deviation of any non-zero asset exists and is non-zero is enough to make V into an inner product space: we define the inner product of two assets using the covariance. It follows from the axiomatic development of Euclidean geometry that you can use theorems about distance in Euclidean geometry to prove results about the standard deviation, and hence the risk, of assets.

One result from Euclidean geometry that is particularly useful is that in an n-dimensional Euclidean space, we can find n orthogonal vectors of length 1. Translating this into the language of risk, we learn that it is possible to find n uncorrelated portfolios of standard deviation 1. Let's call these portfolios $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n$. If we treat these $\mathbf{e}_i$ as the "basic" assets in our market, then any portfolio can be written as a combination $q_1\mathbf{e}_1 + q_2\mathbf{e}_2 + \ldots + q_n\mathbf{e}_n$ and the risk will be given by the length of $|\mathbf{q}|$. In other words: our market is exactly the same as standard Euclidean space, you just need to choose the right coordinates.

4 The two fund theorem

The last section focussed entirely on risk. We should also consider how much a portfolio costs and the mean payoff of a portfolio. Doing this will allow us to prove a the so-called *two fund theorem*.

We'll assume that the cost of a portfolio is just given by the total cost of the component assets scaled by the quantities. This means that the cost, C, is a linear function on the space of portfolios. So we can divide the space of portfolios into slices of equal cost as shown in Figure 3. Each slice is an $(n-1)$-dimensional hyperplane in the space of all portfolios.

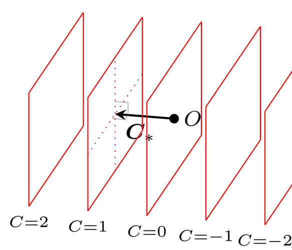

Figure 3: Hyperplanes of constant cost

In Euclidean geometry, given a hyperplane, we can draw a unique line through the origin that intersects the hyperplane at right angles. The line and hyperplane intersect at the closest point on the hyperplane to the origin. This can be proved using Pythagoras's theorem, just as we saw when considering the market of identical stocks.

Translating this into the language of our market, we find a portfolio, $\mathbf{C}_*$, which minimizes the standard deviation among all portfolios of cost $C = 1$. This portfolio is shown in Figure 3.

Similarly, the mean payoff of a portfolio will be a linear function, P, on the space of portfolios. So there is a portfolio, $\mathbf{P}_*$, that minimizes the standard deviation among all portfolios of mean payoff $P = 1$.

Portfolios that have both cost $C = 1$ and mean payoff $P = 1$ will lie in the intersection of the hyperplanes $C = 1$ and $P = 1$. These portfolios form an $(n-2)$-dimensional subspace of the space of all portfolios which we will call L.

The portfolios matching any specified combination of cost and mean payoff will lie in a parallel $(n-2)$-dimensional subspace, L'. By the minimality properties used to define $\mathbf{C}_*$ and $\mathbf{P}_*$, L is perpendicular to the plane OC_*P_* containing the origin, $\mathbf{C}_*$ and $\mathbf{P}_*$. Hence L' is also perpendicular to this plane. The situation is illustrated in the three dimensional case in Figure 4.

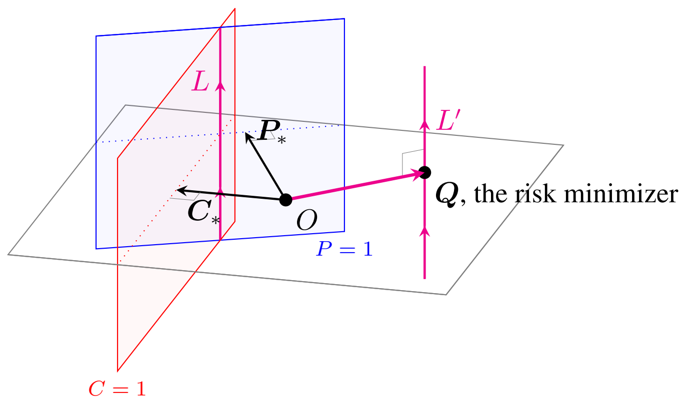

Figure 4: The geometry of the two fund theorem

It follows that if Q is closest point to the origin in the subspace L', then Q lies in the plane OC_*P_*.

Translating this into the language of portfolios we see that the portfolio, $\mathbf{Q}$, of lowest risk among all portfolios, L, with a given cost and given mean payoff lies on the plane OC_*P_*. We have proved the following:

Theorem (Two fund theorem [2])**.** *Given a portfolio* $\mathbf{A}$*, we can find a portfolio,* $\mathbf{Q}$ *with the same cost and mean payoff in the plane* OC_*P_* *that is no riskier than* $\mathbf{A}$*. This portfolio can be written as a linear combination of the two portfolios* $\mathbf{C}_*$ *and* $\mathbf{P}_*$*.*

An investment fund is a financial product that allows you to invest in a portfolio of assets which are selected on your behalf by a fund manager. The fund manager's job is to design funds which will attract investors. The two fund theorem says that a fund manager can cater for any potential customer irrespective of their budget or appetite for taking risks by managing only two distinct funds (assuming investors agree with the fund manager's view on the risk and return of the market). Each customer can then divide their investment between the two funds according to their risk preferences.

Notice that reflection in the plane OC_*P_* preserves the cost, mean payoff and risk of a portfolio. This gives an alternative route to proving the two fund theorem. Any optimal investment must be invariant under these reflections, but the only invariant portfolios lie in the plane OC_*P_*. This symmetry argument lies at the heart of the two fund theorem. It can be applied very generally to financial problems in the Markowtiz market, and to other markets which admit symmetries [author(s)].

5 The efficient frontier

Because of their risk-minimizing properties, portfolios in the plane OC_*P_* are called *efficient* portfolios. We may plot a graph of the risk of each portfolio in the plane OC_*P_*. Since risk is given by distance, the resulting graph will be a cone.

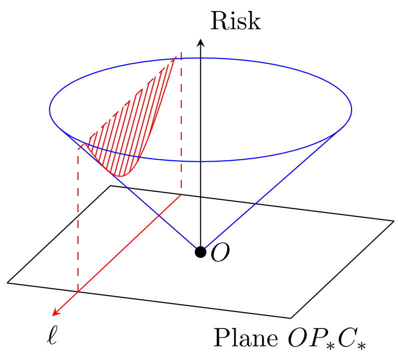

Figure 5: A graph of the risk of each portfolio gives a cone

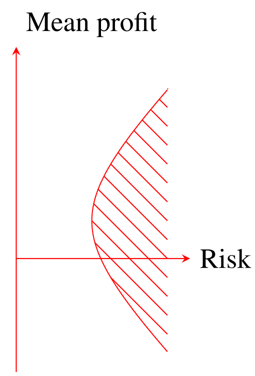

Figure 6: The efficient frontier

If we choose a fixed cost $C = 1$, then we obtain a line of ℓ of efficient portfolios of cost 1. The mean payoff will vary linearly as one moves along ℓ. The risk at each point in ℓ can be computed by intersecting the vertical plane through ℓ with the cone. This is illustrated in Figure 5.

This slice of the cone gives us a graph showing the relationship between mean payoff and risk for efficient portfolios, as illustrated in Figure 6. The hatched area shows all combinations of risk and return that can be obtained in the market. The boundary curve gives the minimum achievable risk for any given mean profit, and is called the *efficient frontier*. The efficient frontier is the iconic image of Markowitz's theory of portfolio optimization. We have proved that the shape of this curve is a conic section, specifically a hyperbola.

6 Conclusion

Euclid's geometry gives an elegant way to understand the relationship between risk and return in financial markets. This demonstrates the power of the axiomatic approach: we may use the same abstract model for what at first appear to be completely different phenomena.

References

[1] H. Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.

[2] R. Merton. An analytic derivation of the efficient portfolio frontier. *Journal of financial and quantitative analysis*, 7(4):1851–1872, 1972.