5

PORTFOLIO OPTIMIZATION

5.1 MAXIMIZING PROFIT, MINIMIZING RISK

A fund manager's task is to achieve two objectives: to maximize profit but at the same time minimize risk. In general, assets that give a higher return bear a higher risk. Government bonds are relatively safe because governments seldom go bankrupt (that happens occasionally). However, the interest that a government bond pays is generally lower than the return on stocks.

The best that a fund manager could do is to base their judgement on their knowledge about the assets available for investment. In other words, they can only do their best to find the expected return and risk for each asset. If their assessments are wrong, then they will make wrong decisions. How to forecast returns and risks is not the subject of this chapter. In this chapter, we shall examine what the fund managers could do based on their assessments.

Ideally, a fund manager should invest all their funds in an asset that pays a higher return and bears a lower risk among all assets in the market. Unfortunately, such an asset does not exist. Even if it does, it will disappear soon. This is because it will be snatched up by participants in the market in no time at all, which means the asset's price will rise, resulting in a lower return to future buyers. If such an asset emerges, and a fund manager spots it and acts

DOI: 10.1201/9781003348474-5

ahead of everyone else, then indeed that is what the fund manager should do.

In general, the fund manager will have to choose among a set of assets of varying returns and risks (according to the fund manager's assessment). They will have to decide how to allocate their funds to individual assets. This is referred to as the "portfolio optimization problem".

When the prices of individual assets in a portfolio change, the fund manager may find that the current portfolio is no longer the best according to their portfolio assessment criteria. The problem of how to adjust the current portfolio is referred to as the "portfolio management problem". To simplify our discussion in this book, we shall not go into this problem (we shall come back to it briefly in Section 5.6). Neither shall we discuss short selling, despite it being important in hedge fund management.

The common saying: "don't put all the eggs in one basket" applies to investment. Diversification is one of the basic principles in portfolio optimization. Dividing one's investments into holdings of two different shares could reduce risk, as long as the two company's share prices do not move up and down at exactly the same rate. If the share prices of the two companies always move in opposite directions, "market risk" (the risk resulted in the price movements in the market) is eliminated. In reality, share prices do not always move in the same or opposite directions. Diversification normally reduces market risk.

5.2 THE MARKOWITZ MODEL FOR PORTFOLIO OPTIMIZATION

The best-known approach to portfolio optimization is the Markowitz model. This model assumes that the fund manager is given a fixed set of available assets to invest in. The fund manager may invest any amount of its capital into any asset. No short selling is allowed. The goal is to maximize return and minimize risk.

Before deciding on how to allocate the funds, the fund manager must assess each asset's expected return and expected risk. The return of an asset may be based on its past return over a given period of time. There are many ways to quantify risk. In the Markowitz model, the risk is measured by the standard deviation of the returns over a period of time. Log returns are commonly used in practice. For daily data, the log return on each day will be calculated.[1] Then the standard deviation of daily log returns will be taken as the asset's risk over that period.

After establishing the return and risk of each asset, one must establish for every pair of assets whether their prices change together in the same direction, in opposite directions or independent of each other. This is measured statistically by the "correlation coefficient". The basic idea is that the more the two assets' movements agree with each other, the larger their correlation coefficient. That means statistically when one asset's price falls, the other's price is likely to fall as well. Holding both assets at the same time is riskier than holding two assets which correlation coefficient has a negative value, which means statistically their prices tend to move in opposite directions.

A portfolio is an allocation of funds into the assets available. For example, for a portfolio with three assets, the allocated funds to them, referred to as "weights", could be 25%, 35% and 40%, respectively. The return of a portfolio is the weighted average of the returns of the assets selected. The risk of a portfolio reflects how much is allocated to each asset, the risk of the individual assets selected as well as how much the prices of the assets move together. Details of the calculation will not be included here. Interested readers are referred to the literature. Readers should bear in mind that the method under the Markowitz model is popular, but not the only way to calculate portfolio risk. In Section 4.4, we described a different way to measure the risk of a portfolio used in the finance industry.

Figure 5.1 plots 10,000 random portfolios for three stocks in the London Stock Exchange from 22 January 2016 to 12 January 2018 (500 days); the three arbitrarily picked stocks are TSCO, BA and RBS.

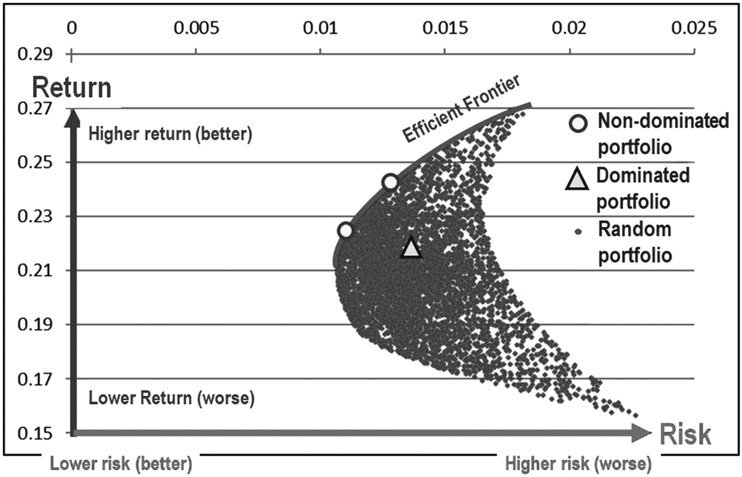

Figure 5.1 The Plotting of 10,000 Randomly Picked Portfolios for TSCO, BA and RBS, Three Stocks from the London Stock Exchange from 22 January 2016 to 12 January 2018 (500 Days) according to the Markowitz Model.

Each portfolio was created by allocating random weights to each of the three stocks. The risks (shown on the x-axis) and returns (y-axis) were calculated according to the Markowitz model.

The fund manager wants to maximize return and minimize risk. In computer science, this is a *two-objective optimization problem*. In Figure 5.1, both portfolios A and B (labelled in circles) are better than portfolio C (labelled in a triangle) because they provide a higher return and a lower risk than C. Portfolio C is said to be "dominated" by portfolios A and B. Instead of holding portfolio C, the portfolio manager will get a higher return and lower risk by holding either A or B. On the other hand, neither A nor B dominates the other, because A has a lower risk than B while B has a higher return than A. Among the randomly generated portfolios, the ones that lie on the top left quarter, labelled "efficient frontier" in Figure 5.1 dominate the portfolios below the frontier.

In practice, fund managers would reduce this problem to a single-objective optimization problem. They might determine the minimum return that it demands and pick the portfolio on the efficient frontier (hence with the lowest risk) that matches that return. Alternatively, they might determine the maximum risk that they are willing to bear and pick the portfolio on the efficient frontier (hence with the highest return) that matches the specified risk.

5.3 CONSTRAINED OPTIMIZATION

The Markowitz model assumes that one can allocate any weight to any asset in the portfolio. This implies that the fund manager can buy a fraction of a share. In reality, shares are bought in "lots", e.g. 100s. Sophisticated investors will be able to get around such constraints, but that complicates the operation and may not be cost-free.

The simplifying assumption impacts the computation: without constraints, the efficient frontier is smooth. That means once the fund manager finds a portfolio on the frontier, it can crawl along the frontier to neighbouring portfolios – crawling can be done by slightly adjusting the weights on some of the assets and assessing the adjustments' impact on risk and return. That way, it can find the portfolio that meets its minimum return or maximum risk requirements. What happens when shares must be bought in whole numbers or in lots? The efficient frontier will become ragged because a small adjustment in weight may not result in a whole number of shares in some stocks. That makes crawling much more difficult because not every point on the efficient frontier contains a valid portfolio. The problem becomes computationally much harder. The task of finding the optimal portfolio (to satisfy either the minimum return or maximum risk requirement) becomes intractable as it is haunted by combinatorial explosion (see Section 2.4).

Buying shares in whole numbers or in lots is just one of the many possible deviations from the Markowitz model. Other constraints may apply. For example, a fund manager may limit the percentage of funds to be invested in the same sector ("don't put all the eggs in one basket" principle). Another fund manager may limit the number of assets held in its portfolio so that it can watch these assets more carefully or limit its transaction costs when it needs to adjust its portfolio.

From a computational point of view, once these constraints are included, the nature of the problem changes further. The techniques best suited for finding portfolios vary depending on the nature of the constraints included. For example, if a large number of constraints are involved, then constraint satisfaction could be employed: The principle of constraint satisfaction is to use constraint propagation to eliminate areas where computation can be saved. In other words, while the constraints are the cause of the complication, they can be deployed to guide the search towards solutions.

In some problems, the fund manager may have some idea of whether certain assets are likely to form good portfolios. Such knowledge could be turned into heuristics in algorithms. Heuristics do not have to be correct all the time. But if they are correct more often than they are wrong, then they could help an algorithm to find better solutions more efficiently. Heuristic search is a very important area of research in AI.

To summarize, adding constraints to the Markowitz model could make the task of finding optimal solutions much harder or even intractable. In this case, knowledge of computation techniques could make a big difference. They could help a fund manager to find better portfolios faster than their competitors. Knowing what to buy or sell ahead of one's competitors is crucial to the success of a fund manager, as explained in Section 1.2.

5.4 TWO-OBJECTIVE OPTIMIZATION

The fund manager has two objectives: to maximize return and to minimize risk. However, in practice (as explained in Section 5.2), fund managers often solve the problem by focusing on one objective. This could be done in many ways, for example:

(1) Determine the maximum risk that the fund manager is willing to accept and find a portfolio with that risk, maximizing return.

(2) Determine the minimum return that the fund manager is willing to accept and find a portfolio with that return, minimizing risk.
(3) Combine the risk and return into one objective; for example, find a portfolio that maximizes the "Sharpe Ratio", which is the return minus the risk-free return divided by the risk.

Many investors would find the question: "how much risk are you willing to take?" difficult to answer. To start, it is difficult for ordinary investors to know how to quantify risks. Besides, the answer to this question depends on what return one is talking about. If the investors want to take as little risk as possible, but the portfolio manager can only find them a portfolio that returns 0.5%, then the investors may be willing to compromise and take on a bit more risk. Similarly, if the investors want to gain a return of 12%, but the best portfolio that the fund manager can find has a 50% chance of losing 90% of the capital, then the investors may be willing to reduce their return expectation accordingly.

Investors in general would find it difficult to determine their maximum risk or minimum return before the portfolio manager finds them a portfolio. In order to specify their preferences, investors need to know the trade-off between risk and return. In other words, to make their decisions, they would benefit from seeing the efficient frontier shown in Figure 5.1.

Given the above analysis, one should be motivated to treat portfolio optimization as a two-objective problem, as opposed to a single-objective problem. Multi-objective optimization is a well-established discipline in computer science. Therefore, it is a bit surprising that multi-objective optimization methods have not found their way into portfolio optimization in the industry.[2] This is probably because portfolio managers do not realize how mature the multi-objective research is and computer scientists with relevant expertise do not realize the opportunity of applying their techniques.

To treat portfolio optimization as a two-objective optimization problem, one does not attempt to find a single solution. Instead, one would attempt to find a set of non-dominating solutions on the efficient frontier. How many portfolios should the portfolio manager present to the investors? That is a question for the investors – they could ask for as many portfolios as they are willing to examine. Instead of determining how much risk to take and how much return to demand in advance, most investors would find it much easier to be given concrete portfolios to choose from. Comparing portfolios is much easier than making abstract decisions. For example, the investors may be given three non-dominating portfolios on the efficient frontier to choose from. Based on the investors' choice, the problem solver may generate two more portfolios on either side of the efficient frontier. This way, the investor's choice can be refined incrementally. Such refinement can be done by mature population-based multi-objective optimization methods.

5.5 THE REALITY IS MUCH MORE COMPLEX

So far, we have explained that the Markowitz model is a simple model for portfolio optimization. With simplifying assumptions, solutions are relatively easy to find. As constraints are added, the model better describes the portfolio manager's true considerations. But finding solutions becomes computationally more demanding. As we treat the portfolio optimization problem as a two-objective optimization problem (instead of reducing it to a single-objective optimization problem), specialized knowledge in computation is required. In this section, we are going to explain that reality is far more complex than what we have described so far.

To start, short selling is not allowed in the Markowitz model. In finance, short selling is an important strategy used by hedge funds. When short selling is considered, risk assessment is more complicated.

Little has been discussed in the literature that in reality, when portfolio managers maximize return, they should maximize return minus cost, where cost should include computation cost and the cost of acquiring expertise. For example, if simulations are used to calculate the portfolio risks (as described in Section 4.4), then fast computers will help to speed up the simulations – we shall come back to the importance of speed later. If the model is complex, expertise is required to compute good solutions and compute them fast. Computational expertise could be expensive to employ.

How much should the portfolio manager pay for computational expertise? In general, the better an expert, the higher it costs, but better solutions could be found fast. However, there is a problem: before employing an optimization expert, the portfolio manager would not know how much they can improve upon the current team. The potential improvement is not even easy to estimate. With so much unknown, it is difficult to determine how much to spend on computation expertise.

Another factor makes the reality even more complex: prices in the market will change and sometimes change fast. That means the optimal portfolio is time dependent. If an algorithm takes too long to compute the optimal portfolio, that portfolio would carry a different return and risk (because the prices of the assets have changed). The algorithm must watch the market while it conducts the computation. It must decide when to return a portfolio that the manager has a chance of acquiring. This is not the same as portfolio management – the problem of adjusting an existing portfolio following the change in asset prices – which is a separate problem.

The real portfolio optimization problem, even without portfolio management, is far more complex than what we have described in previous sections. None of these complications is described in textbooks. The problem is so complex that no computational methods have been developed for it; not even close. A portfolio manager who manages to model better (with as many of the considerations described earlier) and find solutions faster for a more realistic model will have an edge over its competitors. The competition is in research.

5.6 ECONOMICS VS COMPUTER SCIENCE

It is worth elaborating on the point that economists and computer scientists have very different views on portfolio optimization.

A typical computer scientist would say: give me the specification of the problem and I shall find you a solution. This is because, by training, computer scientists typically start with a specification of a problem. Given the specification, a competent computer scientist will pick the relevant techniques for solving the problem.

To the computer scientist, an economist would reply: specifying the problem is the whole of my research! Of course, they are right. The Markowitz model is just a rough approximation of the real problem that needs to be solved. Adding constraints to the specification, as explained above, is just one step closer to the real problem. Solving these approximated problems is not as important as finding a model that captures more of the fund manager's considerations.

Some economists may believe that once they can specify a problem, finding solutions is relatively unimportant. Most economists do realize that the perfect rationality assumption is unrealistic. Some would understand that finding the optimal solution is non-trivial, but they may believe that "if I cannot find the optimal solution, everyone else would have the same problem". Unfortunately, this is not true. Computer scientists would know that for intractable problems, one must settle for suboptimal solutions. Some methods will be able to find better solutions than others and find them faster.

An economist might say: "we are not dealing with a perfect model anyway. What does it matter if two solutions differ by 5% in the quality of their solutions?" The reality is that even if two methods can find solutions of similar quality, speed matters. By using the right algorithms and heuristics, one solver may spend a fraction of the time required by a naïve algorithm to find solutions of the same quality. For example, even the hardest Sudoku puzzle would just take a constraint satisfaction solver one to two seconds to solve on most home computers. A naïve brute force search would be haunted by the combinatorial explosion problem (explained in Section 2.4); it could take years to solve the same problem.

Speed matters especially if the fund manager needs to react to the market. So far, we have left out the problem of portfolio management (the problem of adjusting a portfolio to reflect price changes in the assets). To adjust a portfolio, the fund manager needs to know what the new optimal is. If it takes too long to calculate the new optimal, prices would have changed again. For that reason, computational speed matters.

The economists are correct in pointing out that in portfolio optimization, the research focus should be on the specification of the problem. But not many pay attention to the fact that given a specification, finding solutions is non-trivial. Different computational methods work well for different problems. Depending on the specification of the problem, different methods must be used. Not every computer scientist has the same knowledge about algorithms. Acquiring such knowledge is not free. Therefore, computational knowledge matters.

To summarize, in portfolio optimization, one needs to know what one wants to achieve – that is the task of specifying the problem. One also needs to know how to find good portfolios efficiently – that is the task of algorithms design and selection. To succeed in portfolio optimization, one needs synergy between finance and computation; one needs to know what methods work with what specifications. (Readers may refer back to Section 1.6, when we discussed synergy between finance and computation in trading.)

5.7 SUMMARY

Portfolio optimization is one of the core problems in finance. The fund manager has two objectives to achieve: to maximize return and to minimize risk. Diversification is the basic principle to reduce risk. Many practitioners start with the Markowitz model, for which optimal solutions are not hard to find, especially if one turns the problem into a single-objective optimization problem.

If one wants to better model the fund manager's needs, one must relax the simplifying assumptions and consider realistic constraints. This makes the problem harder to solve. The more factors the fund management considers, the more demanding it is for the fund management team's expertise in algorithms – they need to know what algorithms to use to handle what types of constraints.

There are a lot of opportunities for finance and computing experts to exploit if they could work closer together. The portfolio optimization problem is a two-objective optimization problem. Financial experts probably do not realize how mature multi-objective optimization research is; most computer scientists do not realize the potential of their techniques in this problem. Besides, if they both research deeper into the problem, they will realize that the current models ignore some of the most important aspects of portfolio optimization: computing expertise costs and computation takes time. How much computation expertise should a fund manager acquire? In a market where price change quickly, when should computation terminate to allow the fund manager to acquire a less-than-optimal portfolio? These are all open questions waiting to be studied.

The reality is messy. That is why we need to make models. But we must understand how our models compare with reality.

NOTES

1. The return is the price change in percentage. Log return is often used by industry.
2. The author will not be surprised if some companies are using multi-objective optimization methods in portfolio optimization without publicizing it.