Fundamental Principles of Modern Portfolio Theory & Option Pricing - G D Bernstein

The starting point for modern portfolio theory is the efficient market hypothesis. This comes in three forms but it can be summarized simply in the statement that "it is extremely difficult for the ordinary investor to make money from trading in shares". This is hardly a concept which is likely to upset the typical investment manager. That is until he works out the corollary:- such a high proportion of the money invested in the Stock Exchange is professionally managed that he (the investment manager) may be "an ordinary investor".

The point about all this is not to denigrate the efforts of investment managers, but to point out that the effect of all their efforts will be that, most of the time, shares will be "correctly" priced. That is to say that prices will represent the consensus opinion of the market and nobody will be sufficiently certain that the market price is wrong that they will sell or buy enough shares to affect it. It does not mean that the share price will never change or that the share price correctly foretells the future. After all, there is no reason to suppose that the market will be any more prescient than the people who comprise it.

Having said prices reflect the consensus opinion, that opinion will change from time to time as a result of:

a) News which reaches the market, and
b) Re-assessment of the existing information about the company.

Both of these events are essentially unpredictable. Thus, one would expect share prices to move in a random fashion as they react to new items of information or re-assessments.

There is no reason why some investors should not trade successfully. For example, there is clear evidence that those people who have inside information about companies, for example the Board of Management, do trade successfully (when they are allowed to) and earn above-average returns. One reaction to this has been the establishment of investors' services specifically to monitor Directors' sales and purchases of their own company's shares. However, the lesson for the average user of investment management services is that he needs to choose a manager who is not just very good but substantially above the average of investment managers in the market. His skills will have to be adequate to earn an above average return to cover his fees, the transaction costs he will generate and (we will see below) for the extra risks he takes. It is paradoxical that it is largely a consequence of the investment managers' rapid and discerning response to new information which makes the market so efficient and which largely eliminates the opportunity to profit from that information.

We therefore have reasonably good grounds for supposing, a priori, that changes in stock prices will be fairly random over the short term. David Wilkie's work indicates that there may be a long term auto-regressive element in the UK. (Curiously, the publication of his research may tend to affect the series he studied.) There have been numerous tests of the randomness of the market over relatively short time scales, up to 6 months, which support the notion that prices do move more or less randomly.

The next question is whether all stocks are likely to give the same return in the future. The answer to that question is that on the whole the riskier stocks give a higher return, i.e. the market appears to price stocks in such a way that the riskier ones provide a return which is higher on average than the more staid stocks.

In practice, the definition of risk chosen by MPT practitioners has been "standard deviation of rate of return". This definition has been criticized by UK Actuaries on the basis that it confuses variability with "true risk" [presumably of catastrophe]. In some circumstances, the greatest risk may arise simply from being different from ones competitors - even if one believes that they are wrong! However, if one accepts that price movements are due to unforseeable events, the concept of variability as a measure of risk seems quite sound.

The risk for each share can be split into two parts, "market risk" and "residual risk". Market risk is the part of the variability of a stock's price due to changes in market prices generally. The residual risk is the variability due to the stock's price movements which are uncorrelated with the market. By selecting a diversified portfolio of shares, it is possible to reduce the residual risk to negligible proportions for the entire portfolio. Therefore, one would expect that market prices would only provide a reward for taking on undiversifiable market risk. That is to say that the market will not reward the stock holder for any element of residual risk. At this stage, we are firmly in the realms of theory and it is difficult to test the proposition empirically.

Beta is defined to be a measure of the tendency of a stock to move with the market. Thus a stock with a Beta of 1 would have the same exposure to market movements as the market itself. It might (indeed would) have a substantial residual risk uncorrelated with the market. A stock with a Beta of 0.3 would tend only to move a little in line with the market. An example of low Beta stocks are gold mining shares which tend to be much more influenced by changes in gold prices than by changes in the market prices of shares generally. Thus gold share prices tend to be fairly independent of market movements. Of course, gold mining shares are quite volatile and they have a substantial residual risk which is uncorrelated with market movements.

By combining shares in a portfolio, the residual risk can be diversified away and all we are left with is the market risk of the portfolio. If the shares in the portfolio have an average Beta of less than 1, the portfolio will tend to move rather less than the market as a whole. Conversely, if its Beta is greater than 1 it will tend to move rather more than the market as a whole. It is not automatic for the residual risk to be diversified away. For example, a portfolio which consisted of 100 different gold mining stocks would have a large proportion of the residual risk of its constituents. Nevertheless, for the more usual case of a broadly based portfolio, the residual risk decreases as the number of shares is increased.

The reason for holding shares in the first place is that there is an expectation of earning returns which are greater than from risk-free investments. Thus, we expect that whilst market prices will move up and down, there will be a general upward trend over a period of time. The high-Beta portfolio will tend to magnify the upward trend in prices. Thus, over the long term we could expect to do better with a high-Beta portfolio than with a low-Beta portfolio. Of course, the price that we pay for this is a greater variability in the value of our portfolio over the short term.

Predicting Beta Values

Beta values have usually been calculated by determining the co-variability of a stock with the market over a period of years, i.e. the Beta value is calculated over some period in the past. The past figure is then used as a prediction of the future Beta value for that stock. The process is analogous to attempting to predict the full-grown height of a child by measuring the heights of his parents. In the same way that there is a regression towards the mean amongst humans, there is a similar regression towards 1 for Beta values. Thus, an adjustment needs to be made to the historical figure when using it for prediction purposes.

An alternative approach to predicting the full-grown height of a child is to measure his current weight, physical condition, age etc., and use this information to predict his full grown height. It is also possible to adopt this approach for the prediction of Beta values. Attempts have been made to do this using readily available financial information for companies. Perhaps surprisingly, it is found that the historical method of predicting Beta values works almost as well as the other approach.

It will be apparent from the above that Beta values do vary quite significantly over time. Thus, it is not possible to predict the Beta values of an individual stock with any great accuracy. However, the Beta value of a portfolio is the weighted average of its constituents' Beta values and so the portfolio Beta (for a large enough portfolio) can be accurately predicted. The analogy with predicting the full-grown height of children is irresistible. It is not possible to predict an individual child's height accurately. However, the average full grown height of a class full of children can be predicted with reasonable certainty. In fact, we find that we may need to hold more shares in our portfolio if we are to be able to predict its Beta reasonably accurately then we need simply to reduce the residual risk to reasonable proportions.

Does Variability Matter?

Addressing delegates at the GIRO conference, one of course expects that the answer is yes. However, talking to the trustees or pensions manager of a typical pension fund, one gets the answer that they do not regard risk as an important consideration. This, they say, is because they have a substantial surplus in the pension fund and they do not expect to have to realize investments in the near future. It is worth pointing out to them that the returns on equities have been substantially in excess of the cost (for a large pension fund) of borrowing money. One could therefore ask them why they do not borrow and use the borrowed money to invest in equities, thus retaining the expected turn for the pension fund. The invariable answer to this is that it would be far too risky, by which they mean that nobody else is doing it.

However, it is clear that there is some level of risk, i.e. variability of return, that the pension fund trustees are unwilling to accept. For example, most pension fund trustees would (probably correctly) baulk at the suggestion of borrowing as much money as is in the fund to invest, i.e. gearing up 100%. The same trustees (if one twisted their arms) would probably not be distressed at the suggestion of borrowing 1% of their fund value. There may well be an intermediate level which would be appropriate for the scheme, optimizing the anticipated return without taking on an unacceptable level of risk. Of course, much will depend upon what the pensions schemes in competing companies are doing and who (i.e. the members or the company) will benefit from any excess returns generated or suffer from any losses.

The purpose of the above digression is simply to show, by reductio ad absurdum, that there is some level of risk where the variability of their investments' return does matter to the trustees. They may be able to improve their return if they move their fund's investment stance closer to the point at which they start to feel uncomfortable. At the moment, most fund managers seem to rely on doing much the same as everyone else!

General insurance companies are in the position that they could only invest a part of their assets in equities in any case. The determination of a sound maximum proportion to invest in equities is surely the province of the Actuary. MPT provides him with a tool for analysing the risks inherent in such an investment policy and may enable him to suggest guide lines within which the investment department can work.

Capital Asset Pricing Model (CAPM)

The most diversified portfolio which is available can be constructed by purchasing a small amount of all the stocks in the market place. (The theory runs into some difficulty in defining precisely what we mean by the market place but this is outside the scope of this short paper.) Such a "market portfolio" will obviously have a Beta of 1. The anticipated return on the portfolio will thus be the average return expected from holding shares. The risk inherent in the portfolio will be the risk inherent in holding shares generally.

A conservative investor, who is unwilling to accept this degree of risk, can reduce the risk of his overall holdings by investing a proportion of his money in the market portfolio and the balance in a "risk-free" asset. Thus, if he is only prepared to accept half of the risk of the market portfolio, he would invest only half of his money in it. On the other hand, the more adventurous investor may wish to gear-up his portfolio by borrowing money in order to invest in the market. Thus, any desired degree of risk can be obtained by a combination of investment in the market portfolio plus borrowing cash/investment in the risk-free asset. As an investment philosophy, this approach has not really caught on in the UK. However, as a benchmark for investment performance, it is quite useful. Typically, an investment manager will run an undiversified portfolio, in order to achieve above-average returns. The CAPM indicates that as a benchmark we should not compare his performance with the market but with a combination of the market portfolio plus borrowing or investing of cash, with the same overall degree of risk.

The capital asset pricing model rests on the assumption that there is a generally agreed risk-free asset. A risk-free asset is one which closely matches the investor's overall liabilities. In practice, different investors will have different liabilities. The theory has generally been developed using 3 months deposits as a proxy for the risk-free asset. There is no reason, at all, why index linked bonds could not be used instead if these are considered a better match for the investors' liabilities. At the moment, there is probably little difference between the rates of return on deposits and on index-linked bonds, so this would not make much difference. In the past this has not been the case. However, this, together with problems over the taxation of different investors in the market place, would indicate that the model should be taken as a first approximation rather than as a definitive model of the market place.

Option Pricing Theory

Some definitions:-

A call option : the right to buy shares at a specified price (the "exercise price") at a specified time in the future.

A put option : the right to sell shares at a specified price, etc.

Writer of options: a person who offers one of the above contracts for sale.

A hedged position: a situation where you cannot lose.

Modern option pricing theory is based on "hedging". A hedged position is effectively a matched position. For example, a writer of a call option in a thousand XYZ shares could hedge his position by buying 1000 XYZ shares and an option to "put" those shares at a specified price. Then, if the share price goes up, his call option will be exercised against him but he will be able to deliver the stock. If the share price goes down, he will be able to protect his position by exercising his put option. The existence of this hedge means that there will always be a straight-forward relationship between the prices of a call option and a put option with the same exercise price.

There is a general misconception that the prices of options are more volatile than the prices of the underlying securities. In fact, this is not true. Take as an example, a stock trading at 100p and an option to purchase the stock at 80p. The option price is say 25p. We will ignore both dividends and the interest which can be earned on money for the time being. If the share price reduces by 20p to the exercise price of the option, the option price will also fall. However, the hope value of the option will be more than the 5p appropriate when the stock was "well in the money" at 100p. Thus the option price might fall from 25p to 15p. In this case a 20p fall in the share price has been translated into a 10p fall in the option price. A writer of 1000 call options in the shares could have hedged his position by holding 500 shares. It is this second method of hedging which is the basis of modern option pricing theory.

The reason for the misconception about volatility is that options are always cheaper than the underlying shares. Thus in the above example an investor with £1000 to spend could buy 1000 shares in the company or 4000 options. Naturally, the total value of the 4000 options will move more than the total value of the smaller number of shares.

Clearly, the relationship between the price of the option and the price of the underlying shares is not linear. However, in practice, the curve is relatively smooth and the curvature is not all that great (in most cases). Thus, a writer of call options can hedge his position by purchasing an appropriate number of shares in the underlying stock. the appropriate number of share will depend upon the gradient of the share price/option price curve. If the share price changes significantly, the gradient of the curve will change and the writer of the call option will need to adjust his holding of the underlying shares in the company. Since the gradient generally increases as the share price rises, the writer of the call option will generally have to purchase more shares as the price increases and to sell them as the price decreases.

It is this idea of constructing a hedged position which underlies the well-known option pricing model of Black and Scholes. Their model can be adjusted to take account of dividends, tax, the fact that it is impossible to deal continuously as prices change, and to allow for transaction costs.

Of course the model does not allow for any substantial jumps in share price, for example on the announcement of a take-over. This can move the share price more or less instantaneously to an area where the gradient of the graph is substantially different. There is nothing that can be done about this within the model. As far as I am aware, writers of options have large diversified portfolios of options. There is presumably some ad-hoc allowance in the option price for the possibility of some large jumps in the underlying share price occasionally and they take the rough with the smooth.

It is to be noted that the writer of the call option is not in the least bit interested in the anticipated outlook for the shares, provided he has hedged his position. (He can of course speculate if he wishes to.) His valuation will depend principally upon the volatility of the share price since this will affect the number of occasions when he has to adjust his hedge holding of the underlying shares.

Actuaries familiar with the ideas of immunizing a portfolio of liabilities against a portfolio of fixed interest securities will be aware that, as interest rates change, they have to make adjustments to the underlying portfolio of investments. The notion that the same needs to be done to a hedged portfolio of shares should not, therefore, be particularly difficult.

Selling Short

One of the blithe assumptions of MPT is that it would be straight forward for an investor to borrow money at approximately the same rate that he can deposit it. Also, it is assumed that an investor can sell stock short for an indefinite period. Clearly, for most investors, these assumptions are simply not true and this may give rise to some distrust of MPT.

However, consider a large pension fund which has as part of its normal portfolio an investment of 500,000 British Telecom shares. The fund's natural investment stance is to continue to hold these 500,000 shares.

Suppose that the managers notice a mis-priced option. For example, a BT "put" is too expensive compared to the equivalent "call". They sell 150,000 BT puts and buy 150,000 BT calls but they need to go short of 150,000 BT shares. Simple! They sell part of their existing holding. Since the fund already holds 500,000 shares and would, in the normal course of events, continue to hold those shares, it has in a sense gone short of 150,000 shares if it reduces its holding to 350,000.

Similarly, a fund which would normally expect to hold a cash or short dated bond holding can effectively borrow by running down that cash holding. In this way, relatively small amounts of cash are effectively available to be borrowed at the deposit rate. Larger amounts of cash can be borrowed by substantial investors at rates which are not significantly higher than the normal deposit rates, in any case.

Conclusion

This paper is intended to be a short gallop through the territory of modern portfolio theory. It really can only cover the most basic aspects. Much of the theory can be worked out from the principles set out in this paper, much as the theory of compound interest can be developed from the principles set out in the first one or two chapters of the text book.

No references are given in the paper above since to do the subject justice would require hundreds of references.

Readers interested in studying the subject in somewhat more depth are referred, in the first instance, to the text book on investments by William F. Sharpe (around 650 pages). The definitive text book on options is entitled Options Markets by John C. Cox and Mark Rubinstein (500 pages).