

FROM THE EDITORS

EDITORIAL

In this editorial we shall examine the past for clues as to the expected returns from alternative forms of investment.

If you invest in Treasury Bills or short-term gilts, you receive a fixed rate of interest. This is your compensation for saving money rather than spending it. If you invest in equities rather than Treasury Bills, you are not only giving up the immediate use of your money but you are also taking on risk. Therefore, the difference between the return on equities and the return on Treasury Bills is your reward for this risk. If you do not enjoy taking risk, you will buy equities only if the expected reward for risk is positive. But of course the *actual* reward is rarely the same as you expected. Approximately half the time it will be more and half the time it will be less.

If you invest in long-dated gilts rather than Treasury Bills, you are tying up your money for a longer period. Even if you are content to do this, you still run the risk that unexpected high rates of inflation will erode the real value of your investment. If you do not want to commit your money for a long period or if you do not want to bear the risk of high inflation, you will hold long gilts only if they offer a higher expected return than Treasury Bills.

In order to divide your money sensibly between these different investments, you need some measure of the expected reward for risk from investing in equities and of the expected extra return from investing in long-dated gilts rather than bills. A good place to start is by looking at the performance of bills, equities and gilts in the past.

The Return for Saving

Figure 1 shows the annual return since 1923 from investing in Treasury Bills. This is the compensation that investors have received for saving rather than spending their money.

The average return from bills was 4.2% a year. There were long periods of time when the return was below-average but for each of the past 17 years it has been above-average. There is very little pattern to the changes in these yields and, if all you had to go on was the past record, you could not do much better than assume that the expected bill rate in the future would be the same as the current bill rate.

The Reward for Risk

Figure 2 shows the return on UK equities for each year since 1923. The average of the annual returns was 13.0% but there was considerable variation around this figure. The lowest return in the 58 years was -47.4% in 1974 and the highest return was the 157.1% recovery in the next year.

Approximately 60% of the time the return from equities was higher than the return from bills and approximately 40% of the time it was lower. On average the return from equities was 8.8% higher

than the return from bills. If the past 58 years is representative of the future, this is the reward that you can expect for taking on the risk of equity investment.

Although investors received an average return of 13.0% each year from investing in equities, the *annually compounded* return from equities was less than this. To see why this was so, imagine that equity prices doubled one year and halved the next. In this case the average return was $(+100\% - 50\%) \div 2 = 25\%$. However, prices by the end of the two years were back where they started; so the annually compounded return was 0%. The moral of this story is that, when there is a lot of variability in rates of return, the compound return is much less than the average return in each year. In the case of equities the compound rate of return since 1923 has been 6.1% a year more than the compound rate of return on bills.

The Extra Return from Long Gilts

About nine years in ten the yield on long gilts was greater than the yield on Treasury Bills. The average difference in yields was 1.6%. If on average investors expected bill rates not to change, then this 1.6% is a measure of the extra return that investors expected for giving up the liquidity of short-dated investments.

Of course, in any one year the return from long gilts may be more or less than the promised yield. Figure 3 shows the return from long gilts for each year since 1923. The lowest return was -16.0% in 1974 and the highest was +51.8% in 1977.

Approximately 50% of the time the return from long gilts was higher than the return from bills and approximately 50% of the time it was lower. The average return was 0.7% higher than the return from bills. If on average investors' forecasts of interest rates were correct, then this 0.7% is a measure of the extra return that investors expected for giving up the liquidity of short-dated investments.

Thus we have two, roughly similar estimates of the expected extra return from long gilts. If we look at the past differences in promised yields, our estimate of this extra return is 1.6%; if we look at the past differences in actual returns, the figure is 0.7%.

Inflation

The returns that we have referred to above are measured in terms of money. Unfortunately, the real returns have been much lower than the money return.

Figure 4 shows the annual inflation rate since 1923. The average inflation was 4.8% a year. Thus the *real* pretax returns from investing in Treasury Bills have averaged -0.6% a year, the real returns from equities have averaged 8.2% and the real yields on long-dated gilts have averaged 1.0%.

In the next issue of RMS we shall look at what the past can tell us about the relative risk of these investments.

Richard Brealey Ian Cooper

Fig. 1 TREASURY BILL RETURNS 1923-1980




Fig. 2 EQUITY RETURNS 1923-1980




Fig. 3 LONG GILT RETURNS 1923-1980





Fig. 4 INFLATION RATES 1923-1980



FROM THE EDITORS

Last year the *Institutional Investor* published an article entitled "Is Beta Dead?". Its report of the demise of modern portfolio theory and the rebirth of ignorance included a good deal of sensational journalism. However, the basis for the article was an interesting but controversial paper written several years earlier by Professor Richard Roll of the University of California, Los Angeles. Since the *Institutional Investor's* article attracted considerable interest in Roll's arguments, we decided to devote this editorial to explaining them.

Roll's concerns about modern portfolio theory centre on the fact that the beta of a share measures its sensitivity to movements in a market index. Therefore, our estimate of beta will depend on which market index we decide to use. If we use an inappropriate index, we are likely to misjudge the share's contribution to portfolio risk and we may be led into misjudging the relationship between risk and expected return and making mistakes in building portfolios.

The most common response to this observation has been to suggest that purely negative criticism is of little value unless it is accompanied by constructive alternatives. Modern portfolio theory is almost certainly the best theory that we have on the subject and, until something better comes along, we cannot do more than use the best ideas available.

This response is somewhat extreme. Roll has at the very least provided a useful service in reminding us that, whenever we use simple risk measures, we must try to ensure that they are appropriate for our purposes. Let us, therefore, look at the principal uses of beta in the light of Roll's comments.

Risk Measurement

It is possible to measure the risk of a portfolio by taking into account not only the total risk of the individual shares but also the extent to which each pair of shares moves together. In the same way we can measure the contribution of each share to portfolio risk by calculating the extent to which the share covaries with the portfolio. This would be the beta of the share with respect to that portfolio. Since these procedures do not involve the use of an index, they are immune from Roll's strictures.

In practice, it would be extremely inconvenient if we could never make a statement about a share's risk without having to define exactly what portfolio it was being held in. That is why investment managers find it convenient to measure beta against a standard index and divide risk into two parts — the market risk (or beta) and the specific risk. It is when we make this simple subdivision that we need to choose some measure of the market and, therefore, encounter Roll's criticism that our estimates of beta will depend on the choice of market index.

The problem with Roll's criticism is that the user in this case knows what he is getting: a measure of the

contribution of the stock to the risk of a well-diversified UK equity portfolio. If we use any other well-diversified UK equity portfolio as the benchmark, we will get betas that are similar to those measured against the FTA — All Share Index. The problem does not lie with the risk measures themselves — they are what they are designed to be. The important thing is to remember to use them properly — that is, as a measure of the contribution of the stock to the risk of a well-diversified UK equity portfolio and not as a measure of its contribution to any other portfolio.

Portfolio Selection — Domestic Equities

In the editorial for the last issue of the RMS we described the principles behind active-passive portfolio management. In particular, we saw that the investment manager needs to decide how much to bet on the prospect for individual shares and how much to invest in the passive, or core, portfolio. The precise split should depend on the expected reward and the risk of these alternative investments. This implies that the portfolio manager needs to measure each share's specific risk and beta relative to the passive, or core, portfolio. For most investors the core portfolio will consist of a well-diversified portfolio of UK equities. In that case it is appropriate to measure specific risk and beta relative to the FTA-All Share Index. However, if you are managing a unit trust that is restricted to investing only in oil shares, you should not construct your portfolio on the basis of risk measures relative to the All Share Index.

Portfolio Selection — Wider Portfolios

The contribution of a UK share to the risk of an internationally diversified portfolio cannot be measured directly by its beta relative to the FTA — All Share Index. It can only be measured by its beta relative to an international stock market index.

However, this does not mean that published risk measures are of no interest to the manager of an international portfolio. As long as the prices of UK and overseas equities do not move closely together, we can construct our international portfolio in two stages — first, we construct a sensible UK equity portfolio and then we decide how to divide our money between the UK and overseas. Many firms do in fact appoint a separate manager for the UK portion of the equity portfolio and for the overseas portion; they are effectively using our two stage procedure for portfolio management. The published risk measures for UK shares can be used for the first stage (constructing the UK portfolio) but they cannot help with the second (blending the UK and overseas portfolios).

The Relationship between Risk and Expected Return

Modern portfolio theory states that the difference between the expected return on a share and the interest rate represents the investor's expected reward

for taking on unavoidable market risk. In well-functioning markets the expected reward for risk should increase in proportion to the share's beta *as measured against the market portfolio*.

This theory about the relationship between beta and expected return is generally known as the capital asset pricing model. As the name implies, the capital asset pricing model describes how securities are priced but does not tell the manager anything about what portfolio he should hold. It is quite possible, therefore to disagree with the capital asset pricing model but to believe in the principles of efficient portfolio selection. We stress this because Roll's criticisms are concerned more with the capital asset pricing model than with portfolio selection.

Roll does not criticize the capital asset pricing model as such but only the way that tests of the model have been conducted. He stresses that the model is concerned with the share's beta as measured against the market portfolio. The reason for this is that the model assumes that the market portfolio is efficient in the sense of giving the highest expected return for its level of risk. Therefore, Roll argues, if we want to test the capital asset pricing model, we must test directly whether the market portfolio is efficient.

Roll goes on to point out that tests of the model have usually looked at the relationship between share returns and the betas as measured against some market index. If this index does not give the highest expected return for its level of risk, then the expected reward on each share will be only approximately related to the betas as measured on that index.

There is little doubt that Roll's criticisms of previous tests of the capital asset pricing model have some force. On the other hand, these tests *do* show a relationship between past betas measured against a broad market index and subsequent returns. This finding does not preclude the possibility that betas measured against some other portfolio would give an even closer relationship with returns. However, it does suggest that the capital asset pricing model is at least approximately true. In other words, published risk measures do tell you something about the expected subsequent returns.

Performance Measurement

When measuring a manager's performance it is important to take into account the risks that he has incurred. The usual way to do this is to ask whether the fund would have done better to invest in an equally risky package of short-term debt and the market index. The justification for making this comparison is that, since even an unskilled investor could have chosen this package, a skilled manager should be able to beat it.

Roll's criticism of this technique concerns the choice of the appropriate portfolio for comparison. If the benchmark portfolio is unobtainable (eg because the fund is restricted to stocks with trustee status), then we should not necessarily blame the manager if the benchmark portfolio gives the higher returns. Conversely, if it is not the best passive alternative (eg because the index includes low dividend shares that are unsuitable investments for tax-exempt investors), then we should not necessarily congratulate the manager for beating the benchmark portfolio. Although these problems can be reduced by careful choice of the benchmark portfolio, any comparison with a combination of short-term debt and a market index may not be the best comparison.

Roll's worries about placing too much reliance on measures of portfolio performance are, therefore, certainly valid. However it is important to remember that unusual short-run performance is as often as not the result of good or bad luck, so that you need to examine the portfolio's performance over many years before you can be sure that the result is really due to the manager's skill. The errors that come from misestimating the portfolio's beta are likely to be minor compared with those that come from looking at portfolio performance only over the short term.

None of Roll's criticisms is directed at modern portfolio theory as such or at the ideas behind risk measurement. The real importance of Roll's work is to remind us that we must make sure that the risk measures we use are appropriate for the task at hand.

Richard Brealey Ian Cooper