Yield curves for gilt-edged stocks

A research article prepared in the Bank's Economic Section. The article is largely the work of J. P. Burman and W. R. White (who is now with the Bank of Canada)

1 Introduction

There are fifty-two British government and government-guaranteed securities at present quoted on the stock exchange, ranging from very short-dated stocks with redemption in less than a year to those like Consols and War Loan which are undated. Excluding eleven stocks which are very small issues or have unusual features, such as redemption by drawings or sinking fund, these can be treated as homogeneous; the main feature distinguishing each is its dateof maturity, though its coupon is also important. The yields may be plotted on a chart to examine how they vary with the remaining length of time to maturity. An example is provided in Chart A, where the crosses mark the individual stocks.

The relationship between the maturity of a stock and its relative yield is frequently systematic and broadly obvious from simple inspection; for example, yields on long-dated securities may tend to be above those on short-dated, as in Chart A. In an attempt to capture this systematic relationship more precisely, a yield curve can be fitted to the points on such a chart. These curves are commonly fitted by applying some simple mathematical formula. Thus, for example, curve I in Chart A has been fitted by constraining the shape of the yield curve to be a parabola and then using statistical techniques to discover which one best fits the scatter of individual yields. Fitting a curve, like a parabola, with a low degree (*i.e.* with only a few bends in it) is simple, but it frequently leaves a considerable gap between the fitted curve and some individual yields, so that a large part of the variation between them remains unexplained.

The statistical fit may be improved by using a higher order curve *i.e.* one with more bends. Thus, curve II in Chart A provides a much closer fit than curve I. Indeed, if the curve were allowed to have a sufficient number of bends, it would be possible to obtain one that would pass exactly through each point. Concern with improving the statistical fit, by allowing a more flexible curve, may, however, complicate and even obscure the true relationship between yield and maturity. Thus, the normal process of curve fitting involves a compromise between some low order curve which is simple and informative and a higher order, more flexible curve, which fits better. Yield curves for British government stocks obtained by this kind of curve fitting have been published in the *Bulletin* for the past five years.[1]

Although the form of curve to fit to the scatter of yields is usually chosen on ad hoc statistical grounds, this form nevertheless has implications about the underlying nature and working of the market. Thus in fitting a simple, low order curve, it is implicitly assumed that the market brings about continuous smooth adjustment in the prices and yields of neighbouring stocks throughout the whole range of

1 See the article "Yield curves and representative yields on British government securities" in the March 1967 *Bulletin*, pages 52–6.

Chart A
Illustrative yields on gilt-edged stocks

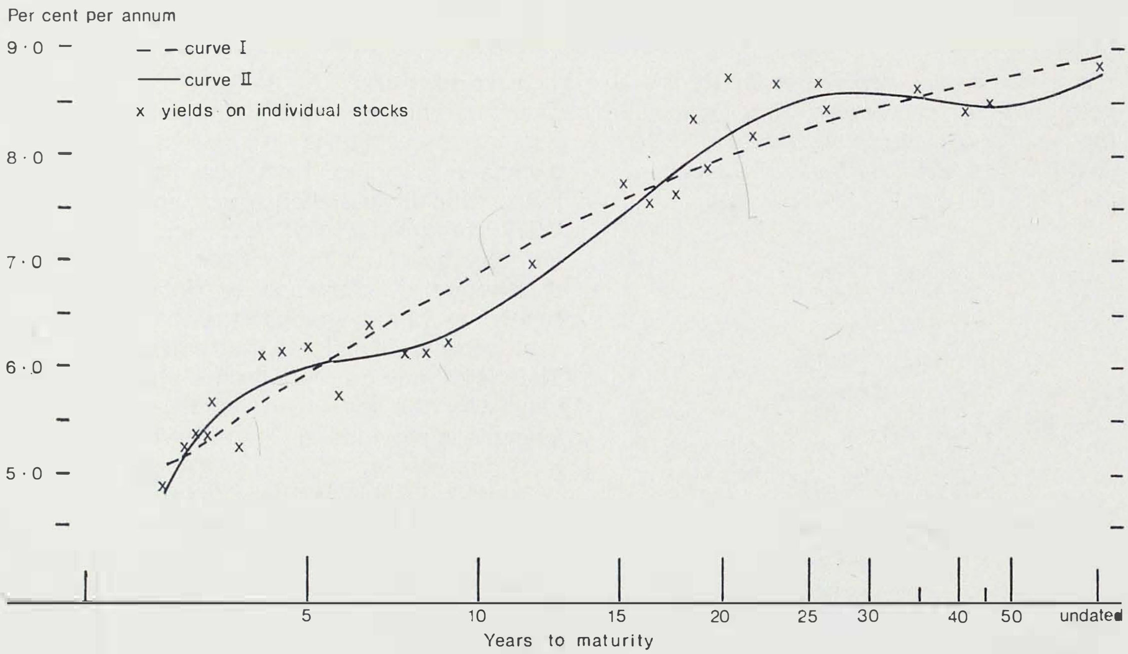

maturities, since the curve has no sharp bends. It is therefore not surprising that academic economists, working with yield curves derived by such curve fitting processes, have tended to provide explanations of how the market functions in terms of these smooth differences in yields, using some version of the expectations hypothesis. This hypothesis holds that the shape of the curve primarily reflects investors' expectations of the future levels of interest rates and prices. On the other hand, many market operators, even though they are often prepared to use yield curves constructed in this fashion, do not believe that the market exhibits such smooth adjustment over all maturities: they regard certain facets of the expectations hypothesis *e.g.* that investors should feel capable of forming clear expectations of relatively far distant events, as implausible; and they often view the market as discontinuous and divided into separate segments. Clearly, if market operators are right, constraining the yield curve to take some simple curvilinear shape is inappropriate, because any subsequent theoretical analysis to explain this shape will inevitably be biased.

This argument suggests that the form of the curve should be derived instead from *a priori* theories about the working of the market and that their validity should be tested by comparing the statistical fit of the resulting curves (see, for example, the discussion of segmentation in Part 2). In short, the form of curve to be fitted should be based on analysis of the working of the market, rather than obtained from a statistical artefact.

The present analysis starts from the view that certain types of investors have different 'preferred habitats'. For example, discount houses deal mainly in stocks under five years from maturity, banks in those below fifteen years, and insurance companies and pension funds with longer-term liabilities tend to concentrate on stocks with lives over fifteen years.[1] Preferred habitats arise from the desire to match the maturity of assets and liabilities, in order to minimise risk. But at the same time investors have expectations about future interest rates: empirical evidence in both the United States and the United Kingdom suggests that these expectations extend for perhaps two or three years ahead. The furthest point at which an investor's expectations are fairly well-formed is called his 'planning horizon' and the time from now till then is his 'decision period'. For some investors the horizon will be very much nearer than their preferred habitat, and an investor may also change his plans before the horizon is reached.

Suppose an investor's horizon is two years and he is holding a stock in his preferred habitat. He has expectations about the yields at his horizon on this stock and on a whole range of alternative stocks. These expected yields imply expected capital gains or losses; and from these in turn, together with the (certain) running yields, expected returns may be calculated. If one of the alternative stocks offers a higher expected return over this two-year period than his present holding, the investor should consider a switch (which could be reversed later). But there is a risk that the actual return on the alternative stock will prove to be lower than expected instead of higher. So the investor can only be tempted away from his preferred habitat if the expected improvement in return from investing in the alternative stock exceeds some minimum amount. This is called his risk premium: the more averse to risk he is, the larger the premium will need to be for a given stock; and the longer the maturity of the alternative stock, the larger the premium will be. This is because a given error in forecasting the future yield will cause a larger error in the future price of a longer-dated stock, and therefore in the expected return over the decision period.

Some investors have no permanent habitat in the gilt-edged market but migrate there temporarily when the climate in other markets seems relatively unfavourable. For them the choice of stock can also be made in terms of expected returns and risk premiums, with horizon equal to the length of their intended investment.

Much of the remainder of this paper, and also the more detailed paper by White [**11**], is concerned with translating this theoretical framework into a more rigorous form which can then be applied to construct yield curves consistent with the underlying hypotheses about the working of the market. This should make it possible to test statistically various alternative hypotheses: for example, whether the market is characterised by continuous switching between stocks throughout the maturity spectrum in response, say, to changing expectations; or whether it is broken up, segmented, into a number of separate compartments with little switching and imperfect substitution between them, perhaps because the risk premiums required by investors to shift from their own preferred habitats are so high.

---

1 For U.S. evidence on preferred habitats, see [**9**].

Furthermore, the construction of a yield curve from an analysis of the working of the market imposes the discipline of having to specify all the factors that might be expected to influence the prices and yields of stocks. Some of these factors, such as the effect of differences in the coupon rate and in the tax position of the various investors in the market, turn out to be quite important. Like many other yield curves, the type published until now in the *Bulletin* did not take explicit account of these systematic factors. (Their importance was implicitly recognised by the exclusion of stocks with coupons less than 5%.) This shortcoming may not be serious when the purpose is to identify stocks whose prices seem out of line, and thus to profit from such apparent anomalies by switching out of, or into, these stocks. Experience will soon reveal that certain stocks, say, those with high coupons, normally stand above the fitted yield curve, and appropriate allowance can be made for these systematic effects.

However, the Bank need a yield curve to assist in judging the appropriate terms for new government issues (which are usually made at a price close to par), and in advising on the rates of interest to be charged by the central government for lending to public corporations and local authorities. For these purposes the shortcomings mentioned above in the traditional type of yield curve are more serious. For example, it happens nowadays that the lower the coupon rate, other things being equal, the lower the yield. The main reason for this is that a greater proportion of the yield to maturity of a low-coupon stock takes the form of a capital gain, which for many investors is free of tax if realised after more than one year, whereas their dividend income is taxable. When interest rates are historically high, the coupon which the Government must offer on an issue to sell at par will also have to be historically high. Existing stocks in the same maturity range with lower coupons will then promise some capital gain on maturity. As a result, a curve statistically fitted to the existing scatter of yields, without an explicit estimate of the coupon effect, will significantly under-estimate the coupon which will have to be offered by the Government on a new issued to be priced at par. (The exclusion of stocks with coupons below 5% only partially rectifies this under-estimation and has the disadvantage of greatly reducing the number of stocks observed.) To meet this criticism, the type of yield curve introduced in this article shows the nominal rate of interest which a stock of each maturity should bear, in order to be issued at a price of 100. This is termed the par yield curve. As explained above, when yields are historically high the par curve lies above most of the yields on existing stocks (see, for example, Charts B and C in Part 3).

In Part 2 below the theoretical framework is elaborated to enable the yield curve to be derived, and in Part 3 the results are summarised. In some respects this exercise has been successful. These curves fit much better than their predecessors, and in particular yield differences due to coupon are very largely explained. Yield curves derived in this manner will therefore replace the previously constructed curves in future issues of the *Bulletin*. Nevertheless some of the results have been disappointing so far. Given the nature of the underlying data it has proved to be very difficult, in some cases virtually impossible, to distinguish between the effects of various factors upon the present structure of yields: namely, the length of horizon, the expected level of future yields, and the risk premium. The only way out of this dilemma at the moment has been to fix certain of these variables at plausible but arbitrary values. But such deficiencies suggest the need to search for further information that will allow additional constraints to be imposed on the variables in the model.

2 Factors determining relative yields on government stocks

The concepts of 'planning horizon' and 'decision period' are familiar to market operators; they form part of the procedures used when advising on switching from one investment to another ([**1**], [**2**], [**7**]). The following example provides a simple illustration of how such concepts may be used.

Suppose an investor feels he can take a reasoned view about the course of interest rates over the next two years but not beyond—his planning horizon is thus two years—and suppose he is considering the choice between investment in a five-year and a ten-year stock. Assume that the two stocks in question have coupons of 6% and present redemption yields of 6%, and that his view about interest rates leads him to expect that at the planning horizon, two years hence, both stocks will have redemption yields of 7%. Their present prices are, of course, 100. Their expected prices in two years' time would, on these assumptions, be 97·34 and 93·95 respectively, for, given the equal changes in redemption yields, the capital loss is greater on the longer-dated stock. Over the decision period of two years the expected returns—coupon payments less the capital loss—are 4·71% and 3·04% per annum respectively. If there were no uncertainty at all surrounding the investor's expectations, he should, in these circumstances, switch his holdings of ten-year stock (if that were his preferred habitat) into the shorter five-year stock, with the intention of reversing the switch at a profit in two years' time. Obviously, expectations are always uncertain to a degree; the effect on his decision of this uncertainty and the risk it generates is considered in the next section.

If all investors' planning horizons and expectations were the same, switching between the two stocks would occur until their prices reflected equal expected returns over the decision period (still ignoring risk). In the above example, if equilibrium were reached with an expected return of, say, 4% on both stocks, the present redemption yields on the five-year and ten-year stocks would then be 5·69% and 6·24% respectively. Thus, the yield curve would slope upwards, in anticipation of the future rise in interest rates to 7%.

This simple illustration can be made more general in two ways. First, instead of choosing between two stocks, each investor will have before him the whole range of stocks maturing beyond his planning horizon, from which he will select the one giving the highest expected return over the decision period. He cannot take into consideration stocks maturing within his planning horizon in the same way—the problems of estimating the short end of the yield curve are discussed later. Secondly, investors are bound to have different planning horizons and different expectations.[7] If, however, the planning horizons of the many investors are distributed over a fairly narrow range—an assumption which as explained in the introduction need not conflict with the idea of a preferred habitat and which will be relaxed later—average expectations may be regarded as relating to an average planning horizon. Thus, in equilibrium, the average expected return over the average decision period should be equal for all stocks, before adjustment for risk.

This framework is still not sufficient to explain the term structure of yields at a given time, unless the expectations of investors tend to follow some rational pattern. Fortunately, there is empirical evidence that most of such expectations relate to a general level of interest rates in the relatively near future (see the studies of the market in the United States by Kane and Malkiel [**4**] and in the United Kingdom by White [**10**]). On this basis it seems plausible to assume that investors, having taken a view about the level of future yields at their relatively short planning horizons, would expect the level to remain unchanged beyond the horizon, in the face of virtually complete uncertainty about developments in the distant future. Another way of putting this is that the horizon is the farthest point at which expectations are definite enough to imply further changes in interest rates. In these circumstances, market expectations about future yields may be represented by a single average figure, the average planning horizon. Thus, in equilibrium the average yield on a stock maturing at the average planning horizon is by definition equal to the expected return over the decision period, it is now possible to construct a testable theory accounting for the structure of yields beyond the planning horizon.

In the absence of risk, the theory would account for upward and downward sloping curves in terms of bearish and bullish expectations respectively. If interest rates are expected to rise, the expected level of yields at the planning horizon is greater than the expected return over the decision period, and the investor requires higher present yields on the longer-dated stocks to offset the larger expected capital losses. If yields are expected to fall, the reverse is true and he is content with lower present yields on the longer maturities, in the expectation of larger capital gains.

The influence of risk aversion

In recent years, yields on the longer-dated stocks have usually exceeded those on the shorter-dated stocks, except in 1960, 1965 and 1967–68. The notion of risk may help to explain this. Indeed, if there were no risk, an investor whose expectations differed from average would put all his funds into the one stock which to him offered the highest expected return, after taking account of transaction costs.

---

7 Lintner [**5**] has argued that with a common horizon but differing expectations prices should reach an equilibrium determined by the weighted average of these expectations, where the weights would be proportional to each investor's financial resources.

The main reason that economists offer to explain the desire to diversify portfolios is risk aversion. As investors clearly cannot forecast with certainty what yields will be at their planning horizons, they are inevitably uncertain about the expected return over their decision period. The effect of this uncertainty increases with the length of the period to maturity—the longer the stock, the greater the capital loss resulting from a given error in forecasting future yields; and the effect of uncertainty is also greater for a low-coupon stock than for a high-coupon stock of the same life. In equilibrium the risk-free return expected over the decision period is the (certain) yield on a stock maturing at the planning horizon. But for stocks maturing beyond this horizon, the theoretical framework is modified to take account of risk by assuming that the expected return over the decision period must exceed the (certain) yield by the amount of a risk premium which increases with the period to maturity.

A specific measure of the risk premium is needed and alternatives are discussed in White [**11**, section IIA]. The measure used here is the percentage change in the price of a stock caused by a 1% change in the yield. It is applied to the expected price, and is formally defined as the ratio of a very small percentage change in price to the corresponding change in yield. It is proportional to the measure of volatility of a stock as calculated by brokers; and is familiar to economists as the interest-elasticity of the price. The volatility is zero when a stock matures at the horizon and normally reaches a maximum for an irredeemable stock; but, since it varies inversely with the coupon, the volatility can actually exceed this 'maximum' for a very long-dated low-coupon stock standing well below par.

The influence of taxation

A fundamental feature of the market for British government securities is the distinction which persons and most companies draw between the investment return obtained as income (dividends), which is liable to tax, and the return received as capital gains, which is tax free if the stock is held for more than one year.[1] Investors for whom this distinction is important are called net investors. Other investors for whom dividends and capital gains are equivalent are described as gross investors. They are mainly non-profit-making organisations, which pay no tax on either form of investment return, and authorised dealers in securities including jobbers, banks and discount houses, which pay the same rate of tax on both.

These tax considerations give rise to differences in gross yields between high and low-coupon stocks, the latter having greater appeal for net investors because they offer a larger capital gain at redemption. If the market were dominated by gross investors, gross redemption yields of high and low-coupon stocks with the same maturity would tend to be equal. If the market were dominated by net investors (paying tax at, say, the standard rate), the net redemption yields calculated at this rate of tax would tend to be equal. In practice the situation is somewhere in

1 Except between April 1965 and April 1969.

between, with gross redemption yields of high-coupon stocks above those for similar low-coupon stocks, but not enough above to equalise the net yields. This suggests that the equilibrium prices are a weighted average of two sets of prices: one set derived from equal gross returns over the decision period and one from equal net returns over the decision period. The simplest way of incorporating this tax effect in the theoretical model was found to be the creation of an 'effective tax rate'; this is used to calculate an effective yield – representing the average of the net return for net investors and the gross return for gross investors.[1]

There are two other factors which will affect differently high and low-coupon stocks. In the first place, the risk premium on low-coupon stocks should be larger because of their greater price volatility. This means that the prices of these stocks will tend to be lower (the yields higher) and the yield gap between high and low-coupon stocks smaller than tax effects alone would suggest. In the second place, expectations about future shifts in interest rate levels will influence prices (and yields) of high and low-coupon stocks differently. If prices are expected to rise (yields to fall), low-coupon stocks will be more attractive because of the greater prospective capital gain to be obtained, so that their present price will rise (present yield fall) relative to the price of high-coupon stocks with the same period to maturity. Conversely, when prices are expected to fall, the prospect of capital losses will make low-coupon stocks relatively less attractive. In sum, the 'coupon gap' will tend to widen in a bull market and to narrow in a bear market. But the evidence shows that these two factors are dominated by the tax effects.

The basic theoretical model summarised in the previous section can now be reformulated more realistically. First, even if a flat yield curve at and beyond the planning horizon can be assumed, it would seem necessary to allow for the permanence of yield differences associated with the coupon. Thus, it is postulated that investors expect high-coupon securities to continue to have higher gross redemption yields than low-coupon securities of the same maturity. This postulate is introduced into the model by defining a flat expected effective yield curve at the planning horizon, on the assumption that effective yields of individual stocks (*i.e.* at the effective tax rate) are equal at this horizon.

Secondly, as explained above, the expected return over the decision period (which investors as a whole are assumed to be trying to equalise for all stocks) is an average of returns to gross and to net investors, calculated with tax at the effective rate. Although the model determines equilibrium prices in terms of these effective yields, the Bank are concerned, as described above, with fitting a par yield curve, which requires the calculation of gross redemption yields from these prices.

---

1 Evidence from the stock registers about the relative importance of gross and net investors is too vague to provide an estimate of the effective tax rate, mainly because of the large amount of nominee holdings. The effective tax rate, expressed as a percentage of the standard rate of tax, was derived from the process of fitting the yield curve: see Appendix 2.

Segmentation between different parts of the market

The theoretical framework developed so far generates smoothly upward or downward sloping curves, accounting at the same time for consistent differences between the yields on low and high-coupon stocks. It does so because it is supposed that all investors consider the total range of stocks maturing beyond their planning horizons, notwithstanding their preferred habitats. However, such a theory remains incapable of explaining the occasional humps and troughs which are sometimes evident in the pattern of yields. Indeed, the evidence reported by White [**11**] covering the period 1961 to 1971 suggests that a theory which assumes all British government securities to be substitutable one for another provides an inadequate explanation of the actual term structure of yields.[1] Imperfect substitutability implies that there may be significant differences between planning horizons, a possibility quite familiar to market operators but assumed away in the development of the theory so far. If, however, such differences exist, they will mean that investors may be broadly divided into groups according to their preferences to operate in particular segments of the market. In these circumstances, the model developed thus far would only be valid within each segment, because expectations bearing on different segments will refer to different time periods in the future and so are almost bound to be different.

The model was, therefore, extended to allow first for two distinct markets and secondly for three markets. In the event, allowance for the possibility of three segments gave unsatisfactory results (see Appendix 1) so that preference was given to a model which allowed for two distinct markets, in short-dated stocks on the one hand and medium/long-dated stocks on the other. The theory was then applied separately to each of these segments of the market. Thus, each group of investors, according to the segment of the market in which it operates, is assumed to have its own average planning horizon and decision period, with associated expectations about future yields and returns. For each group, the effective tax rates are also likely to differ, because of the varying importance of gross and net investors; and, in principle at least, the risk premium could also differ as between the two groups. As it turned out, it was not possible to estimate the size of the risk premium for either group with any precision, so it was arbitrarily decided to assume that the premium increased at the same rate in each segment according to the period from the planning horizon to maturity.

It seems most unlikely that investors' preferences would give rise to two sharply distinct segments in the market, for there will always be some investors prepared to operate in both and so contemplate switching between the two segments; and in any event the statistical identification of the different planning horizons associated with each segment cannot be very precise. As the observed shape of the yield curve tends to be fairly smooth, it seems more consistent to suppose that the relative importance of investors

---

1 This remains true even if stocks with sinking funds or other special features are excluded.

with shorter and longer planning horizons changes gradually with maturity, so that the final curve can be constructed by splicing the two separate curves together. In the case where the shorter planning horizon is one year and the longer three years, the average expectations of investors operating in the short segment will determine a smooth yield curve starting from the one-year point, whereas the yield curve relating to the long segment will begin at the three-year point. For stocks maturing more than three years hence, the curve generated by the short segment of the market is given gradually decreasing weight and the curve associated with the long segment increasing weight in the composition of the final curve. In practice, a smooth combined curve can be obtained with a splicing band of four years or more.

To summarise, the theoretical model accounting for the observed term structure of yields defines prices and so gross redemption yields on all British government securities in terms of the following nine parameters:

(1) and (2) Short and long planning horizons.

(3) and (4) Expected level of net yields at the short and long horizons.

(5) and (6) Risk-free expected net returns over the decision periods up to the short and long horizons.

(7) and (8) Effective tax rates for the two segments of the market.

(9) Maximum risk premium (*i.e.* the rate relevant to irredeemable stocks), which determines risk premiums at all maturities in both segments.

The formulae underlying the model are in Appendix 2.

3 Results

The estimation of the par yield curve proceeded in two steps. First, the values of these nine parameters were varied to find the combination which gave rise to estimated gross redemption yields most nearly approximating to actual gross redemption yields—in other words, the values which provided the best statistical fit. Secondly, from this best set of parameters it was possible to estimate the par yield curve, a continuous curve tracing out the gross yield on a hypothetical stock standing at 100 (net of accrued interest). The definition of the par yield curve is explained more fully in Appendix 1.

As it happens, not every one of the nine parameters can be estimated unambiguously; in an analysis confined to the yields on different stocks at one point of time, the theoretical model is over-specified. For example, when the length of either of the planning horizons was varied, while holding fixed the values taken by the other parameters, the statistical fit of the theoretical curve changed significantly. However, when both the long planning horizon and the corresponding expected return were varied together, theoretical curves of almost equally good fit could be obtained; and this was also true for the short planning horizon. While some experts suggest that the most likely lengths of the short and long planning horizons are of the order of six months and two to three years respectively, more plausible values for the expected returns over the decision periods were obtained with somewhat longer planning horizons. As a compromise, the planning horizons were arbitrarily assigned values of one year and four years, for it was recognised that reliable estimates of the true average planning horizons cannot be ascertained from the kind of evidence used.

A rather more serious problem which was encountered in the process of estimation was the virtual impossibility of assigning a value to the risk premium. This difficulty stems from the fact that the data used record the results of decisions to buy or sell stocks and so embody indistinguishably the many factors such as expected returns and risk premiums which may have led to the decisions being taken. For example, the observed steepness of the curve for the long segment of the market can be attributed either to bearish expectations or to strong risk aversion. To complete the estimation of the theoretical model of yield curves, it was necessary to select and hold fixed a value for the maximum risk premium.

This dilemma regarding the risk premium can be further illustrated. Throughout 1971 prices were rising (yields falling) and extremely large amounts of stock were sold by the authorities. Even so, the pattern of yields remained sloping upwards more steeply than usual. Taking the maximum risk premium as 1%—that is, over the four-year decision period the expected return from an irredeemable stock would be 1% higher than the certain return from a four-year stock—implies that expected yields at the planning horizon were higher than actual yields at that time, whereas common sense suggests that in such a bull market the reverse would be true. To account plausibly for the readiness of investors to purchase such a large volume of stocks, and mainly long-dated ones at that, it is necessary to put the risk premium as high as 3% to 4%, if not more. In these circumstances, the yield gap between high and low-coupon stocks opened up, indicating an increase in the relative importance of net investors who prefer low-coupon stocks, or perhaps resulting from the greater supply of high-coupon stocks.

With such an apparently large risk premium, some movement of investors out of their preferred habitats might have been expected. It is not easy to find direct evidence of this; but, for example, during this period building societies bought a substantial amount of stock maturing beyond fifteen years, even though they normally hold short and medium-dated stocks. This suggests a readiness on the part of investors to consider, at least to a limited extent, stocks which do not necessarily match their liabilities, provided the promise of a higher yield is sufficient.

Finally, it was found that experiments with the width of the splicing band to secure a single curve representing the pattern of yields suggested that the most satisfactory band was one of four years.

The data used to test the model related to end-quarters from June 1969 (just after capital gains tax had been removed) to September 1972. Estimates of the parameters of the model and tests of goodness of fit are given in

Chart B

British government securities: curves of gross redemption yields at 30.3.72

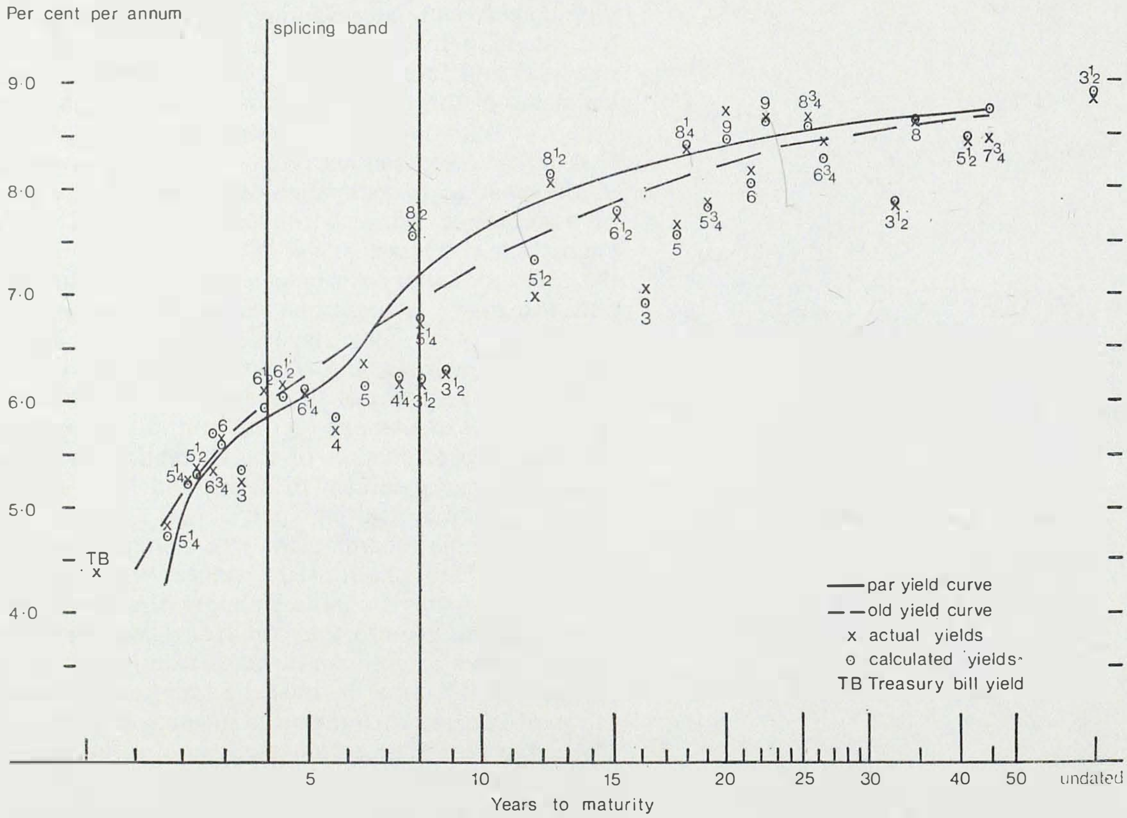

Chart C

British government securities: curves of gross redemption yields at 30.6.72

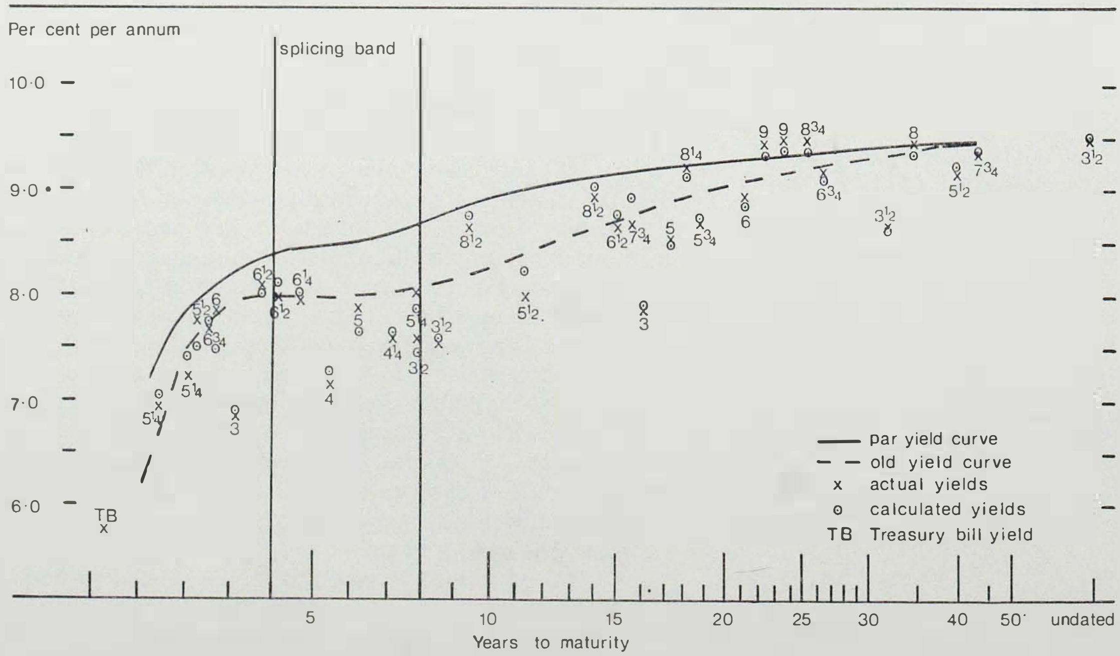

Appendix 2. Charts B and C show the actual and calculated yields on individual stocks at 30th March and 30th June 1972; and the goodness of fit is measured by the differences between them. Only one irredeemable stock, $3\frac{1}{2}\%$ War Loan, is included in the data and the curve is terminated at the last dated stock, because it cannot be accurately estimated beyond this point. The upper line traces out the par yield curve, while the lower represents the yield curve derived on the old basis. In fitting the old curve, Treasury bills and stocks maturing in more than six months were included in the data. For the par curve, stocks below the horizon (one year) cannot be included, as the theory does not predict their yields—unless another, still shorter, horizon were introduced (see below). Indeed the curve becomes unreliable for times shorter than the shortest stock included in the data, so it has not been extrapolated beyond this point.

One reason for this unreliability is that small price movements of very short-dated stocks correspond to large yield movements. A more important reason, probably, is that the model takes no account of the yield on Treasury bills, although this becomes progressively more influential on the yield of a short-dated stock, given the knowledge that the Issue Department of the Bank of England is willing to buy it in due course at a price related to the Treasury bill yield when it becomes the next maturity. Furthermore, such a stock with a life of less than eighteen months is especially attractive to banks, since it becomes a reserve asset when its life drops below one year. These reasons suggest that stocks maturing within two years may form a separate segment of the market with a very short horizon—and with lower expected yields when the yield curve slopes upwards. But it is impracticable to subdivide the short-dated market further, because there are too few stocks to permit estimation of a separate portion of the curve.

Chart D compares the par yield curves for recent high and low points of the market and the last working day in October.

In conclusion, it must be admitted that the theoretical model which has been developed has not been fully tested and so its validity still remains open to some doubt. This is particularly so with regard to the risk premium. Different values were assigned to the risk premium but the effect of these on the final curve was negligible, which implies that the estimates of the expected yields and returns over the decision period cannot be considered as wholly reliable. None the less, this approach to the estimation of yield curves, which first develops a theory accounting for the behaviour of investors in the marketfor British government securities and subsequently tests it against the observed data, would seem to hold more promise than the more common practice of fitting yield curves using arbitrary mathematical formulae. In particular, the new approach largely takes account of differences caused by the size of coupon.

Chart D
Time/yield curves of British government securities[a]

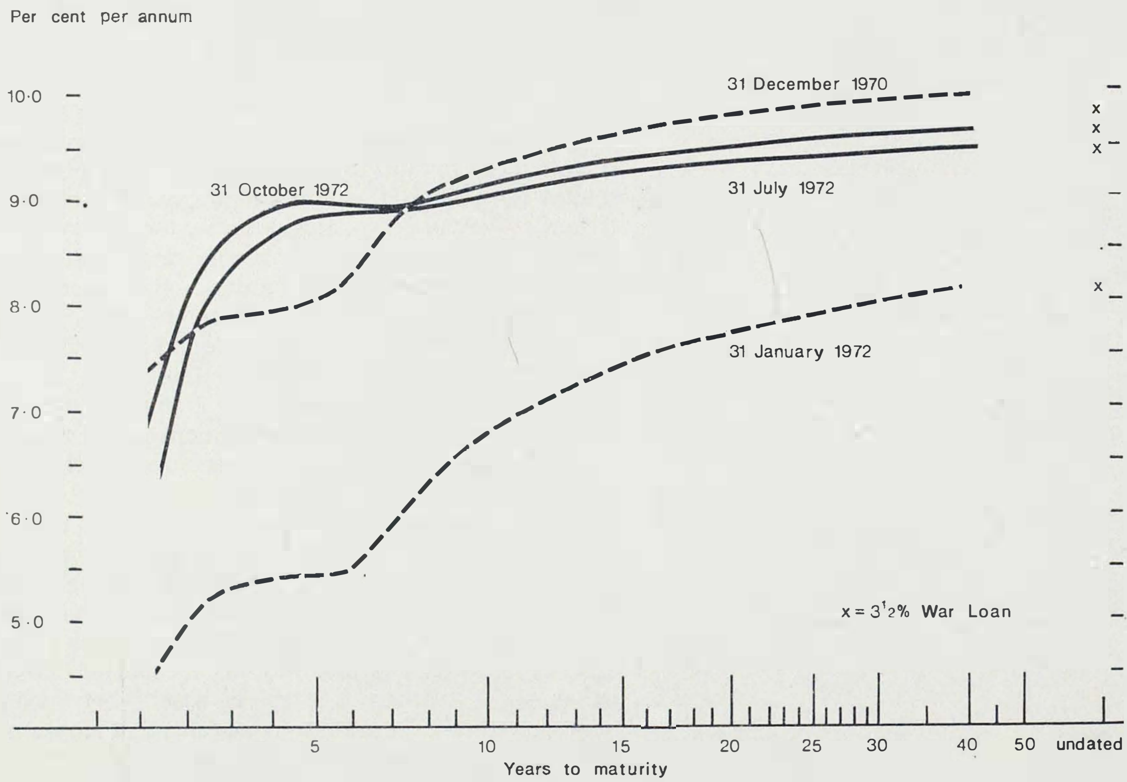

a The lines have been fitted to the gross redemption yields. The curve runs from the shortest-dated stock with a life of more than one year to the longest-dated stock.

References

1 **Brew, J. M.** "Gilt-edged yield curves" *The Investment Analyst,* No. 16 December 1966 pages 3–22

2 **Grant, A. T.** "Switching of British government securities" *The Investment Analyst,* No. 3 August 1962 pages 14–29

3 **Jones, Alcwyn** "Spiral—a new algorithm for non-linear parameter estimation using least squares" *The Computer Journal,* Volume 13 August 1970 pages 301–8

4 **Kane, E. J. and Malkiel, B. G.** "The term structure of interest rates: an analysis of a survey of interest-rate expectations" *The Review of Economics and Statistics,* Volume XLIX No. 3 August 1967 pages 343–55

5 **Lintner, John** "Security prices, risk and maximal gains from diversification" *The Journal of Finance,* Volume XX No. 4 December 1965 pages 587–615

6 **Marquardt, D. W.** "An algorithm for least squares estimation of non-linear parameters" *Journal of the Society for Industrial and Applied Mathematics,* Volume 11 No. 2 June 1963 pages 431–41

7 **Pepper, G. T.** "The selection and maintenance of a gilt-edged portfolio" *Journal of the Institute of Actuaries,* Volume 90 Part II No. 385 1964 pages 63–103

8 **Powell, M. J. D.** "An efficient method for finding the minimum of a function of several variables without calculating derivatives" *The Computer Journal,* Volume 7 No. 2 July 1964 pages 155–62

9 **Terrell, W. T. and Frazer, W. J., Jn.** "Interest rates, portfolio behaviour and marketable government securities" *The Journal of Finance,* Volume XXVII No. 1 March 1972 pages 1–35

10 **White, W. R.** "Expectations, investment and the U.K. gilt-edged market—some evidence from market participants" *The Manchester School of Economics and Social Studies,* Volume XXXIX No. 4 December 1971 pages 293–314

11 **White, W. R.** "The term structure of interest rates—a cross-section test of a mean variance model" *Issues in monetary economics* H. G. Johnson and A. R. Nobay, editors Oxford, forthcoming

Appendix 1

Accrued interest and the par curve

A word needs to be said about accrued interest and its effect on the par curve. The quoted price of a stock (except those maturing within five years) drops immediately on going ex-dividend; it rises with the accrual of interest through the dividend date, continuing until the next time the stock goes ex-dividend. If the market contained only gross investors, the rise and the fall would be equal to the gross dividend; but, as has been seen, the data indicated a balance between gross and net investors, leading to the concept of an effective tax rate. The model implies that the price should rise between ex-dividend dates by the amount of dividend net of tax at the effective rate. But the par curve is based on gross yields and the price rises more slowly than is needed to keep the gross yield constant. Therefore the gross yield rises as accrued interest increases and falls back abruptly at the ex-dividend date. This is the reason why gross investors prefer to buy 'dirty' (*i.e.* with maximum accrued interest) and sell 'clean' (*i.e.* with negative accrued interest); and of course net investors have the opposite preference.

For short-dated stocks the quoted price excludes accrued interest, but the influence of net investors tends to affect the total cash payment (including gross accrued interest) in exactly the same way as for longer-dated securities.

The par yield curve may be thought of as derived from hypothetical stocks existing at every maturity date. Since redemption always occurs at a dividend date, the price of such a hypothetical stock sold by the Issue Department of the Bank will usually contain some accrued interest. However, if the hypothetical stocks are defined as new issues, there is no accrued interest. (If a stock's initial life is not an exact number of half-years, the first dividend is always reduced *pro rata* to the number of days from the date of issue.) Therefore the par yield curve is defined as measuring "the nominal rate of interest which a stock should bear in order to be issued at a price of 100 without accrued interest".

Segment boundaries

The most generally satisfactory curves were obtained with horizons of one and four years. The curves generated in the two segments were spliced together over the band four to eight years. If the number of segments is increased to three, simple inspection and market knowledge suggest another boundary at around fifteen years. Experiment shows that in order to avoid abrupt kinks, the second splicing band needs to be rather wider than that between the first two segments, about six years. The band from twelve to eighteen years was chosen.

The results with three segments showed some improvement in goodness of fit of the curve. But the shapes showed several bends and were largely determined by the position and width of the splicing bands. Also, the parameter values became extremely unstable, indicating that the segments (especially the middle one) contained too few stocks to permit estimation of so many variables. It was concluded that the two-segment model was more satisfactory.

Appendix 2

1 Single horizon (risk and tax ignored)

Data: t = life of stock (in years)
r = coupon/100 (payable half-yearly)
P = actual price of stock/100

Parameters: t_o = decision period (in years)
Y = expected horizon yield/100
x = expected decision-period return/100

Derived variables: $t' = t - t_o$
$\hat{P}$ = calculated price of stock/100
P_e = expected price of stock at horizon/100
$\hat{y}$ = calculated redemption yield/100

First calculate P_e (value at the horizon of future flow of dividends and redemption money):

$$P_e = \frac{\frac{1}{2}r}{1+\frac{1}{2}Y} + \frac{\frac{1}{2}r}{(1+\frac{1}{2}Y)^2} \dots + \frac{\frac{1}{2}r+1}{(1+\frac{1}{2}Y)^{2t'}}$$

i.e.

$$P_e = \frac{r}{Y}[1-(1+\tfrac{1}{2}Y)^{-2t'}] + (1+\tfrac{1}{2}Y)^{-2t'}. \quad (1)$$

[In this formula and throughout the appendix it is assumed that t and t_o are exact multiples of half-years. The formulae actually used are slightly more complicated when this is not so.†]

In the same way calculate $\hat{P}$:

$$\hat{P} = \frac{\frac{1}{2}r}{1+\frac{1}{2}x} + \frac{\frac{1}{2}r}{(1+\frac{1}{2}x)^2} \dots + \frac{\frac{1}{2}r+P_e}{(1+\frac{1}{2}x)^{2t_o}}$$

i.e.

$$\hat{P} = \frac{r}{x}[1-(1+\tfrac{1}{2}x)^{-2t_o}] + (1+\tfrac{1}{2}x)^{-2t_o}P_e. \quad (2)$$

Then $\hat{y}$ is obtained from $\hat{P}$ by solving:

$$\hat{P} = \frac{r}{\hat{y}}[1-(1+\tfrac{1}{2}\hat{y})^{-2t}] + (1+\tfrac{1}{2}\hat{y})^{-2t}. \quad (3)$$

2 Single horizon, risk included

The risk function or volatility is defined as:

$$V = -\frac{Y}{P_e}\frac{dP_e}{dY}$$

using P_e as in (1).

This becomes:

$$V(t') = \frac{\frac{r}{Y}[1-(1+\frac{1}{2}Y)^{-2t'}] + (Y-r)t'(1+\frac{1}{2}Y)^{-2t'-1}}{\frac{r}{Y}[1-(1+\frac{1}{2}Y)^{-2t'}] + (1+\frac{1}{2}Y)^{-2t'}}. \quad (4)$$

Note that $V(0)=0$ and $V(\infty)=1$.

Define: x = expected risk-free decision-period return/100
$x(t')$ = expected decision-period return/100, allowing for risk
and $x(t') = x + aV(t')$. (5)

The calculated yields are then obtained from equations (1), (2) [with $x(t')$ substituted for x], (3), (4) and (5). If $a = 0{\cdot}01$, the maximum risk premium is 1%.

† The redemption yield formula in this case was kindly provided by D. T. Harris of Mullens & Co.

3 Single horizon, risk and tax included

α = standard rate of income tax

k = fraction of standard rate which gives effective tax rate in market for government securities.

All yields x, $x(t')$, Y are now interpreted as effective net yields (*i.e.* net of tax at the effective rate). Equations (1), (2), (4) and (5) remain true provided r is replaced throughout by the net coupon $= (1-\alpha k)r$. Equation (3) remains unchanged because the calculated yields are still gross. The parameters x, Y and a are grossed up *i.e.* divided by $(1-\alpha k)$, before being printed out in the table of results, as they are easier to interpret like this.

4 Segmented market

Parameters: t_s = *short* horizon $[t_s \leqslant 1]$

Y_s = expected *short* horizon yield (flat curve)

x_s = expected risk-free *short* holding-period return

t_l = *long* horizon

Y_l = expected *long* horizon yield (flat curve)

x_l = expected risk-free *long* holding-period return

k_s, k_l = fractions of standard rate of income tax relevant to *short* segment or *long* segment respectively

a = maximum risk premium

t_b = mid-point of splicing band between short and long segments *i.e.* band runs from t_l to $(2t_b - t_l)$ $[1 < t_l < t_b]$

$t' = t - t_s$

$t'' = t - t_l$.

There are now two versions of equations (1), (2), (4) and (5), one for each segment. All yields and returns are to be interpreted as effective net yields.

If β is the rate of capital gains tax, then when $t_s < 1$, P_e in (2) is replaced by $\hat{P} + (1-\beta k_s)(P_e - \hat{P})$ and the new equation solved for $\hat{P}$. Similarly if $t' < 1$ or $t'' < 1$, the last term in equation (1) is multiplied by $P_\epsilon + (1-\beta k)(1-P_e)$, where k is k_s or k_l respectively. The new equation is then solved for P_e.

For $t \leqslant t_l$, the equations for the short segment are applied to obtain $\hat{y}$. For $t \geqslant 2t_b - t_l$, the equations for the long segment are applied. For $t_l < t < 2t_b - t_l$, values of $\hat{y}$ are calculated from both sets of equations and combined using the weight function:

$$w(t) = \tfrac{1}{4}(t^3 - 3t + 2) \qquad (6)$$

where

$$t = t_{sp} = (t - t_b)/(t_b - t_l).$$

Thus $w(-1) = 1$, $w(1) = 0$, and the function has zero derivatives at these two points. Yields are defined over the splicing band as:

$$\hat{y}(t) = w(t_{sp})\hat{y}_s(t) + [1 - w(t_{sp})]\hat{y}_l(t). \qquad (7)$$

5 Least squares

The function to be minimised is $\Sigma\ (y_i - \hat{y}_i)^2$ where suffix i refers to individual stocks and $\hat{y}$ is obtained from equations (3) and (7), given values of the parameters. There are various methods of minimising a non-linear function. In the earlier part of the work reported here Powell [8], which minimises a general function, was used. But more efficient methods exist where the function takes the form of a sum of squares – notably Marquardt [6] and Jones [3]. The method now in use is based on Jones, but differs from it in some respects: upper and lower limits are set for each variable, which act as reflecting barriers.

If a variable persistently moves towards a limit, it is fixed and the search continued in a smaller number of dimensions. This ensures rapid convergence.

Arguments can be put forward that the random errors in the model should be price deviations. $\Sigma\ (P_i/\hat{P}_i - 1)^2$ was minimised for some of the data, which should have the effect of giving greater weight to the observations at the longer end. However, the fit at this end was not improved and this alternative was not pursued further.

6 Par yield curve

This is obtained from equations (1) and (2) by setting $\hat{P} = 1$ in (3), so that $\hat{y} = r$, where r is the coupon of a hypothetical stock and so a function of t. Thus (2) becomes:

$$1 = \frac{r}{x}[1 - A] + AP_c$$

where

$$A = (1 + \tfrac{1}{2}x)^{-2t_0}$$

and equation (1) is unchanged.

Solving these:

$$r(t) = Y\frac{1 - A(1 + \frac{1}{2}Y)^{-2t'}}{B - A(1 + \frac{1}{2}Y)^{-2t'}}$$

where

$$B = Y(1 - A)/x + A. \qquad (8)$$

This is a smooth monotonic curve starting at $t = t_0$, $r(t_0) = x$, whose derivatives are also monotonic.

7 Parameters and goodness of fit

Because of extreme multi-collinearity between Y_l, k_l, and a, it was necessary to fix one of them in order to obtain convergence. Experiment showed that it was best to fix a, as the correlation between Y_l and k_l is the lowest, though still over 0·9. This means that in 1971, for example, when external evidence indicated a bull market, if a is fixed at 1%, Y_l and k_l are probably both over-estimated.

It can be easily shown that there is multi-collinearity between x_l and t_l and between x_s and t_s. Replacing half-yearly payment of interest by continuous payment in (8) means that powers are replaced by exponentials:

$$r(t) = Y\frac{1 - \exp(-xt_0 - Yt')}{B - \exp(-xt_0 - Yt')}.$$

The exponent is:

$$-xt_0 - Yt' = -Yt + (Y - x)t_0$$

and

$$B = Y[1 - \exp(-xt_0)]/x + \exp(-xt_0)$$

and if t_0 is small

$$\simeq Yt_0 + 1 - xt_0$$
$$= 1 + (Y - x)t_0.$$

Thus $r(t)$ only depends on x and t_0 through the quantity $(Y - x)t_0$, if t_0 is small, which in practice it is.

So it was decided to fix t_s and t_l at one year and four years respectively, because these gave the most plausible values of x_s and x_l.

The parameters estimated by the model are shown in Table A (with standard errors beneath).

The first section of Table A shows the estimated values of the expected interest rates (net rates as defined in the model, but grossed up at the effective rate of tax, for convenience). The second section shows the estimated effective tax rates: the rapid growth in these rates indicates that gross yields on high and low-coupon stocks have been moving further apart since the abolition of long-term gains tax on government securities.

Table A

Key

y_{tb} = yield on Treasury bills (for comparison).

x_s, x_l = short and long horizon expected decision-period returns (%) ignoring risk.

Y_s, Y_l = short and long horizon expected yields (%).

Fixed parameters: short horizon=1 year, long horizon=4 years.

Figures in italics are standard errors.

(1) Expected interest rates

		Maximum risk premium $=a$											
		0				1%				2%			
Last working days	y_{tb}	x_s	Y_s	x_l	Y_l	x_s	Y_s	x_l	Y_l	x_s	Y_s	x_l	Y_l
1969 June	8·12	10·11	9·47	8·66	9·62	10·11	9·38	8·67	9·21	10·12	9·30	8·68	8·82
		·37	*·37*	*·35*	*·49*	*·37*	*·36*	*·35*	*·45*	*·37*	*·36*	*·35*	*·42*
Sept.	8·05	10·73	9·68	8·90	9·28	10·75	9·60	8·91	8·87	10·76	9·52	8·91	8·49
		·91	*·44*	*·42*	*·60*	*·91*	*·43*	*·42*	*·56*	*·91*	*·43*	*·42*	*·52*
Dec.	7·88	9·81	9·28	9·14	9·01	9·82	9·21	9·14	8·63	9·83	9·14	9·15	8·27
		·66	*·41*	*·37*	*·51*	*·66*	*·40*	*·37*	*·48*	*·66*	*·40*	*·37*	*·45*
1970 Mar.	7·38	8·54	7·94	9·07	8·60	8·55	7·87	9·08	8·27	8·56	7·81	9·08	7·95
		·51	*·32*	*·28*	*·35*	*·51*	*·32*	*·28*	*·32*	*·51*	*·32*	*·28*	*·31*
June	7·04	8·08	7·41	7·64	10·40	8·09	7·35	7·65	9·99	8·09	7·29	7·67	9·60
		·50	*·40*	*·35*	*·48*	*·50*	*·40*	*·35*	*·45*	*·51*	*·39*	*·35*	*·42*
Sept.	6·99	7·08	7·23	7·52	10·15	7·09	7·17	7·51	9·78	7·10	7·10	7·53	9·42
		·43	*·31*	*·26*	*·35*	*·43*	*·30*	*·26*	*·33*	*·43*	*·30*	*·27*	*·31*
Dec.	7·00	7·33	8·20	7·72	10·88	7·33	8·15	7·73	10·48	7·34	8·08	7·75	10·10
		·62	*·48*	*·37*	*·52*	*·62*	*·47*	*·37*	*·48*	*·62*	*·47*	*·37*	*·46*
1971 Mar.	6·79	6·34	7·53	7·07	10·04	6·35	7·47	7·07	9·71	6·36	7·41	7·08	9·41
		·44	*·42*	*·37*	*·48*	*·44*	*·41*	*·37*	*·45*	*·45*	*·41*	*·37*	*·43*
June	5·71	5·24	7·22	6·36	10·45	5·24	7·16	6·37	10·12	5·24	7·10	6·38	9·81
		·44	*·43*	*·34*	*·45*	*·44*	*·42*	*·34*	*·42*	*·44*	*·42*	*·34*	*·40*
Sept.	4·84	4·02	7·09	5·62	9·72	4·03	7·03	5·63	9·44	4·03	6·97	5·64	9·18
		·45	*·40*	*·33*	*·38*	*·45*	*·40*	*·33*	*·36*	*·45*	*·39*	*·33*	*·35*
Dec.	4·48	3·64	6·45	4·77	9·48	3·65	6·40	4·78	9·21	3·65	6·35	4·79	8·96
		·43	*·48*	*·38*	*·46*	*·43*	*·47*	*·38*	*·44*	*·43*	*·46*	*·38*	*·42*
1972 Mar.	4·39	2·95	6·97	5·00	9·83	2·95	6·91	5·01	9·55	2·96	6·85	5·02	9·28
		·46	*·34*	*·23*	*·29*	*·45*	*·33*	*·23*	*·27*	*·46*	*·33*	*·23*	*·26*
June	5·76	6·71	8·96	7·66	10·03	6·72	8·92	7·67	9·71	6·73	8·85	7·68	9·41
		·53	*·44*	*·25*	*·34*	*·53*	*·44*	*·25*	*·32*	*·53*	*·43*	*·25*	*·31*
Sept.	6·79	5·78	9·59	8·00	10·10	5·78	9·51	8·00	9·77	5·79	9·44	8·01	9·45
		·48	*·39*	*·28*	*·38*	*·48*	*·38*	*·28*	*·35*	*·48*	*·37*	*·28*	*·33*

The expected return over the decision period on an irredeemable stock is $(x_l + a)$.

(2) Effective tax rates (as % of standard rate of income tax)

	Maximum risk premium $=a$					
	0		1%		2%	
Last working days	Short	Long	Short	Long	Short	Long
1969 June	26·9	15·0	27·0	13·9	27·1	12·8
Sept.	44·7	8·3	44·9	7·2	45·1	6·0
Dec.	52·0	15·3	52·2	14·4	52·4	13·5
1970 Mar.	58·0	30·9	58·2	30·2	58·3	29·5
June	54·7	34·6	54·8	33·7	54·9	33·0
Sept.	45·4	44·8	45·4	44·0	45·3	43·4
Dec.	69·3	50·0	69·5	49·2	69·5	48·4
1971 Mar.	51·8	72·1	51·9	71·6	51·9	71·1
June	52·5	78·4	52·6	77·9	52·6	77·4
Sept.	52·8	95·6	52·7	95·3	52·6	95·0
Dec.	47·7	95·6	47·8	95·5	47·7	94·9
1972 Mar.	38·9	90·0	38·8	89·6	38·8	89·3
June	73·6	75·5	75·4	75·0	75·6	74·4
Sept.	48·3	67·8	48·4	67·2	48·6	66·6

Table B

Last working days	0 RMS error(%)	0 $\bar{R}^2$	1% RMS error(%)	1% $\bar{R}^2$	2% RMS error(%)	2% $\bar{R}^2$
	Maximum risk premium $= a$					
1969 June	·214	·455	·215	·451	·216	·445
Sept.	·262	·526	·263	·524	·264	·520
Dec.	·236	·452	·237	·449	·238	·443
1970 Mar.	·173	·884	·173	·883	·175	·881
June	·212	·950	·213	·949	·214	·949
Sept.	·160	·974	·161	·973	·162	·973
Dec.	·223	·958	·223	·957	·224	·957
1971 Mar.	·220	·952	·220	·952	·221	·952
June	·207	·974	·207	·974	·207	·974
Sept.	·194	·973	·194	·973	·195	·973
Dec.	·226	·970	·226	·970	·226	·970
1972 Mar.	·143	·987	·143	·987	·144	·987
June	·158	·962	·157	·962	·157	·963
Sept.	·180	·954	·179	·954	·178	·955

Table B gives the root mean square (RMS) errors and the multiple correlation ($\bar{R}^2$) (both corrected for degrees of freedom). The goodness of fit improved remarkably between mid-1969 and mid-1970. The main reason was that the yield curves became steeper, so that the total variance increased, while the unexplained variance (the RMS errors) remained fairly stable at around $\frac{1}{4}$%.

As remarked above, the size of the risk premium had practically no effect on the goodness of fit, using either the RMS error or the multiple correlation as a criterion. The calculated yield curves agreed almost everywhere to within 0·01%. A zero risk premium gave very slightly better results in all cases, but, because risk-aversion is believed to influence behaviour, it has been decided to adopt a conventional value of 1% for this parameter.

The estimates of x_s are below the Treasury bill yield for nearly all dates since March 1971 : reasons why the former may be somewhat unrealistic are given at the end of Part 3 of the article.

The higher effective tax rates for the long-dated segment ($k_l > k_s$) in the latter half of the data period suggest that this segment contains a greater proportion of net investors than the short-dated one. This seems to be consistent with some partial evidence from the stock registers.