Estimating market interest rate and inflation expectations from the prices of UK government bonds

By Mark Deacon and Andrew Derry.(1)

Market expectations of interest rates and inflation can give important insights into the credibility of monetary policy. For that reason, the Bank has carried out extensive research over the past two years into the ways in which inferences about these expectations can be drawn from the market prices of government bonds. The preliminary results of this work were discussed at a one-day conference in March organised by the Bank, in which researchers in the field from other central banks and a number of academics took part. Their comments—and the results of further research in the Bank—are reflected in this article. The article explains the important issues of estimation and interpretation which arise in this work, and outlines possible changes to the techniques the Bank uses in its analysis of market expectations—for example, in its Inflation Report*. Comments on the issues raised and on the methodological changes proposed will be most welcome.*

Introduction

The disruption in the foreign exchange markets during 1992 and 1993 led to the suspension of sterling's membership of the exchange rate mechanism (ERM) and to the widening of the fluctuation bands of most of the remaining currencies. The move from essentially a fixed (though adjustable) exchange rate environment to one of floating currencies has meant that the Bank now relies more heavily on monetary policy indicators other than the exchange rate when assessing monetary conditions. These include indicators of inflationary pressures, inflation expectations and perceptions of the monetary policy stance. For instance, estimates of market interest rate expectations can provide an insight into whether participants expect interest rates to rise or fall in the future. The Bank reports on these indicators in its *Inflation Report*. In addition, knowledge of market inflation expectations can be a useful input into decisions about the funding of the public sector borrowing requirement.

The yield curve obtained from government bond prices has long been used as a source of information about *interest rate expectations*: both its level and slope are useful monetary policy indicators. More recently, however, emphasis has focused increasingly on the implied forward rate curve. This contains the same information as the yield curve, but presents it in a way that allows expectations of interest rates in the short, medium and long term to be distinguished more easily.

Information on *inflation expectations* has until recently typically been obtained from surveys. But these have shown themselves to be unreliable—perhaps because survey respondents have little incentive to answer accurately. There are other drawbacks to using survey evidence: surveys take time to compile, and so may not give accurate estimates of current inflation expectations; and they usually survey only short-run expectations.

Some efforts have also been made to infer financial markets' inflation expectations from asset prices. For example, between 1985 and 1987 the New York Coffee, Sugar and Cocoa Exchange traded futures contracts on the US Consumer Price Index. Using the prices of these contracts, it was possible to obtain a direct estimate of inflation expectations. But no similar contracts on the UK Retail Price Index (RPI) have ever existed, so other means of deriving inflation expectations from the prices of financial assets must be sought. Research has shown that generally only the prices of assets with fixed nominal rates of return contain accessible information on inflation expectations; and that only in the case of government bond prices is it likely that such information can be extracted satisfactorily. This article describes how, by comparing the yields on conventional and index-linked bonds, a measure of inflation expectations can be obtained. It also examines the ways in which such estimates of interest rate and inflation expectations may differ from 'true' expectations.

Deriving interest rate expectations from gilt prices

When pricing financial instruments, participants in financial markets are—either explicitly or implicitly—revealing information about the interest rates that they believe are appropriate for the transactions they are making. So it is possible to gain an insight into their interest rate expectations by calculating the yield or internal rate of return on the instruments. These yields will also reflect other factors—such as the liquidity of the securities, the effects of taxation and the perceived risk of default by the issuer. But the underlying interest rate (for a given level of risk) should

(1) This article is based on two working papers written while the authors were in the Bank's Economics Division: Deacon, M P and Derry, A J, 'Deriving Estimates of Inflation Expectations from the Prices of UK Government Bonds', *Bank of England Working Paper No 23*, and Deacon, M P and Derry, A J, 'Estimating the Term Structure of Interest Rates', *Bank of England Working Paper No 24*.

Some terminology

The gilt-edged market currently consists of around 70 different bonds, the majority of which are *conventionals*. These entitle the purchaser to a stream of cash flows consisting of regular (semi-annual) fixed interest—or *coupon*—payments, and a redemption payment together with the final coupon payment on the gilt's maturity date.

Chart A
Yield curves(a)

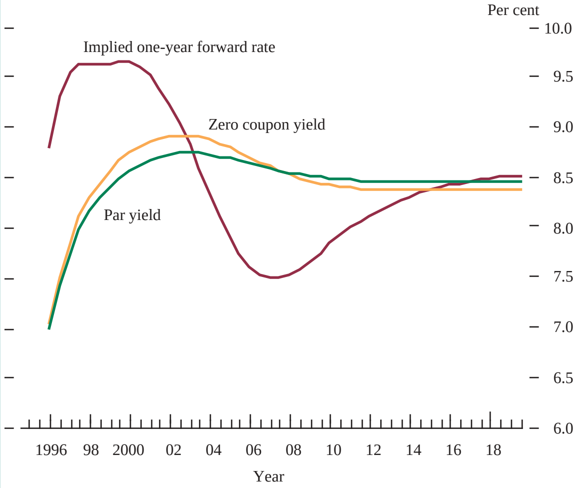

(a) Based on prices on 27 June 1994.

The most commonly used measure of a bond's return is the *gross redemption yield*—the single rate that, if used to value each of the bond's cash flows individually, equates the bond's total value to its price. But two bonds with the same maturity and different coupons may not have the same yield, since the composition of their returns is different—the higher-coupon bond provides more of its return in coupon payments than does the lower-coupon bond. Since (other things being equal) investors prefer assets that provide a return sooner, they are willing to pay a premium for high-coupon bonds. This effect makes it difficult to interpret gross redemption yields—and any measures constructed using them—so other measures have been developed to avoid these interpretation problems. The most fundamental of these is the rate at which an individual cash flow on some future date is discounted to determine its value today—the *spot interest rate* or *zero coupon yield*. It can be thought of as the yield to maturity of a (hypothetical) zero coupon bond, and as such is an average of the single-period rates out to that maturity. The term structure of spot rates, or *zero coupon yield curve*, is the curve which is usually referred to when talking about the term structure of interest rates.

Chart B
Discount function(a)

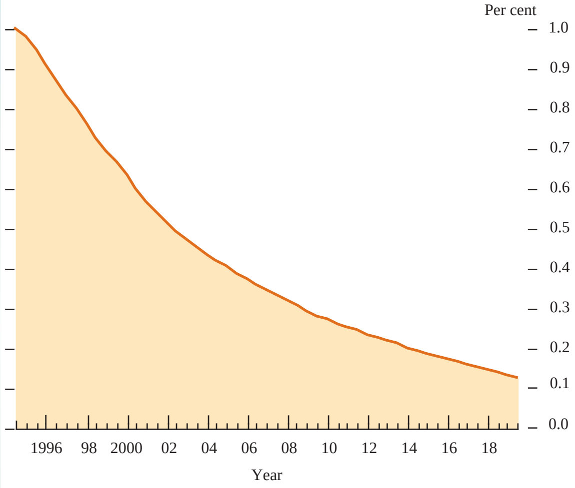

(a) Based on prices on 27 June 1994.

The zero coupon yield curve can be transformed uniquely into three other useful curves: the par yield curve, the discount function and the implied forward rate curve. The *par yield curve* shows the coupons that bonds would require in order to trade at their face value—at 'par'. The *discount function* is a continuous function of discount factors—the value of the discount function for any maturity t is the value today of £1 repayable in t years. Finally, the *implied forward rate curve* consists of implied future one-period interest rates—that is, the one-period rates expected to obtain at future dates. It contains the same information as the spot rate curve but, because it is in effect a marginal curve (whereas the spot rate curve gives an average of expected rates over the chosen horizon), it presents it in a way that makes it easier to interpret for monetary policy purposes. Charts A and B give examples of the four curves.

in principle be unique for each maturity; so, when trying to recover underlying interest rates from market prices, the aim is to construct a curve expressing interest rates as a unique function of time to maturity—the *term structure of interest rates*.

Domestic currency government securities are generally used in the estimation of the term structure of interest rates, since they are normally regarded as being free of default risk. If there were single-payment, liquid government bonds maturing at every future date, the interest rates on them could be used directly to construct the term structure. In the United Kingdom (as in other countries), however, government bonds—*gilt-edged securities*—are not equally spaced through the maturity spectrum: there are many 'gaps' in which one needs to interpolate in order to construct a continuous term structure. Moreover, there are no single-payment—*zero coupon*—gilts, so the technical task of identifying the underlying term structure is further complicated by the existence of periodic interest payments.

Estimating the curve

The first problem in yield curve estimation is how to fill the gaps in the maturity spectrum. A key decision to be taken concerns the shapes that the term structure should be allowed to take—in other words, what trade-off to make between the ‘smoothness’ of the curve (removing ‘noise’, such as pricing anomalies, from the data) and its ‘responsiveness’ (its flexibility to accommodate a genuine movement in the term structure). The purpose to which the term structure is to be put is clearly relevant to this decision. For monetary policy analysis, there is less need than when pricing financial instruments for a precise fitting of local anomalies; a method of estimation better able to generate a smooth curve is preferable.

The Bank’s recent research has investigated a number of different models for estimating the term structure—those due to McCulloch, Schaefer, Nelson and Siegel, and Svensson (an extended Nelson and Siegel model)—as alternatives to the current Bank of England model.(1)

The *Bank of England* yield curve model estimates a par yield curve, essentially by fitting a curve through redemption yields so as to minimise the sum of squared differences between the observed and the fitted yields. The functional form used for the yield curve is known as a *cubic spline* and can be thought of as a number of separate cubic functions joined ‘smoothly’ at so-called knot points.

The other four models are all fitted to a discount function, an approach pioneered by McCulloch. The choice of the functional form in these cases reflects not only the choice between smoothness and responsiveness, but also the fact that a discount function must conform to certain prerequisites based on economic theory—in particular, it should be both positive and ‘monotonic non-increasing’,(2) and be such that the present value of £1 receivable today is £1.

Like the Bank model, the standard *McCulloch* model uses a cubic spline as the functional form.(3) *Nelson and Siegel* start from a different perspective, by specifying a simple functional form for the *forward* rate curve. From this, it is straightforward to derive equations for the term structure of interest rates and the discount function; again, it is the discount function that is fitted by the estimation procedure. An important property of this model is that it is constrained to produce asymptotically flat forward rates for long maturities—a property shared by the Bank model, because of the type of spline used. *Svensson* increases the flexibility of the original Nelson and Siegel model by adding two further parameters, though he considers the standard model to be generally satisfactory for monetary policy applications.

Tax effects

A second major consideration in deciding how best to estimate the yield curve is the choice of method to model tax effects. Tax rules can materially affect the prices of bonds and, if their effects are ignored in the modelling process, can distort the estimate of the term structure of interest rates.

A substantial proportion of investors in the gilt market are taxed at their marginal rate on any coupon income they receive, but are exempt from taxation on capital gains. Bonds with high coupons provide more of their return in the form of coupon income than do low-coupon bonds; so investors who face a non-zero marginal income tax rate but who are not taxed on capital gains will—other things being equal—prefer low-coupon to high-coupon bonds. This preference on the part of tax-paying investors will increase the prices of low-coupon bonds relative to those of high-coupon bonds, a distortion that needs to be removed when attempting to measure the underlying term structure.

The McCulloch method for modelling tax effects consists in estimating a single ‘effective’ tax rate for all maturities. In contrast, Schaefer argues that there is no unique term structure of interest rates, but rather a series of tax-specific term structures, each of which should be estimated using only those bonds which are ‘efficiently’ held by investors in that tax bracket. Schaefer’s specification of the problem highlights a number of difficulties with McCulloch’s approach. First, McCulloch’s effective tax rate will be some kind of average of all income tax rates faced by investors, rather than the marginal rate of the investor whose trading choices determine bond prices. Second, this tax rate is (implicitly) assumed to apply to all bonds along the length of the curve, which is unrealistic if any category of investors has preferences about the maturity of debt held.

The Bank of England model tackles tax effects by explicitly modelling the relationship between yield and coupon, as well as that between yield and maturity. It does this by using capital-income curves which describe the trade-off between capital gain (assuming the bond is held to maturity) and income. Using this method, it is possible to estimate the tax rate faced by the category of investors who determine the price of each bond. Both the Nelson and Siegel and the Svensson models ignore tax effects.

Although Schaefer’s approach of producing tax-specific term structures is well suited to an individual or an institution facing a known marginal tax rate, it has drawbacks for estimating a single ‘market’ term structure of interest rates. The Bank method, though theoretically less rigorous, has distinct practical advantages in this respect and so remains the Bank’s preferred approach.

Other issues

Since the choice of model for tax effects is independent of the curve-fitting technique, the choice between fitting a par yield curve or using a discount function is largely independent of the tax model. It is to some extent a matter

(1) The models are presented respectively in: McCulloch, J H, ‘The tax-adjusted yield curve’, *Journal of Finance*, 1975; Schaefer, S M, ‘Measuring a tax-specific term structure of interest rates in the market for British government securities’, *The Economic Journal*, 1981; Nelson, C R and Siegel, A F, ‘Parsimonious modelling of yield curves’, *Journal of Business*, 1987; Svensson, L E O, ‘Estimating and interpreting forward interest rates: Sweden 1992–93—first draft’, Institute for International Economic Studies, Stockholm University.
(2) A function f is said to be a *monotonic non-increasing function* of time if $f(t_2)$ ″ $f(t_1)$ for all times t_1 and t_2 such that $t_1 < t_2$.
(3) Schaefer (and others) have criticised the specification of a cubic spline on computational grounds—it can introduce significant rounding errors. Schaefer’s model instead uses a linear combination of Bernstein polynomials, which give better approximations to the derivatives—which is important since the forward curve depends on the first derivative of the discount function. Other research has suggested the use of ‘B-splines’.

of taste and beliefs about market behaviour. The discount function approach is explicitly consistent with economic theory, but can be very difficult to estimate; the resultant forward rate curve is also sensitive to small changes in the discount function. The approach of fitting a par yield curve, although theoretically less attractive, appears more robust in practice (particularly when producing implied forward rate curves); this may indicate that it better reflects market pricing realities.

There is also a choice to be made between minimising yield errors and price errors, which can produce significant differences in the forward rate curve. The appropriate choice will again depend on the purpose behind the estimations and—following from that—on the maturity range in which greatest precision is desired. Since the focus in monetary policy analysis is on interest rates rather than prices, it makes sense to minimise yield rather than price errors. A further argument for minimising yield errors is that it improves the fit of the curve at shorter maturities[(1)]—the most interesting from a monetary policy perspective.

Results of comparative testing

To choose between the different curve-fitting approaches, the Bank's research involved carrying out comparative tests of the Bank, McCulloch, Nelson and Siegel, and Svensson models—with each adjusted to incorporate the Bank tax model to ensure fairness of comparison.[(2)] Charts 1 and 2 show the different forward rate curves for two recent dates (the Nelson and Siegel curve is excluded because of its closeness to the Svensson curve for these particular dates). They illustrate the sensitivity of the forward rate curve to the choice of fitting approach, and in particular to the constraints set on the long end.

The Bank's provisional judgment is that the economic constraints imposed by the discount function approach—

Chart 1
Implied forward rates[(a)]

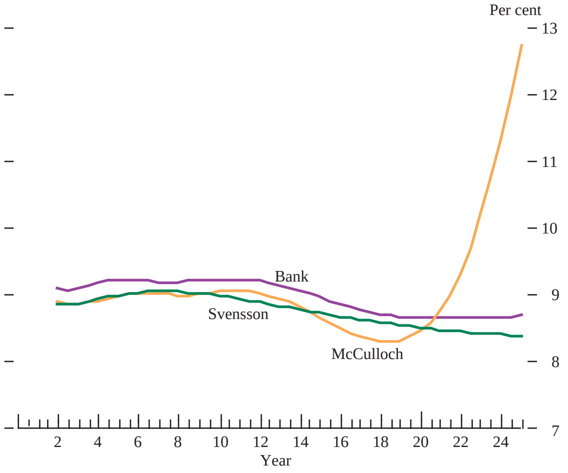

(a) Based on prices on 22 June 1992.

Chart 2
Implied forward rates[(a)]

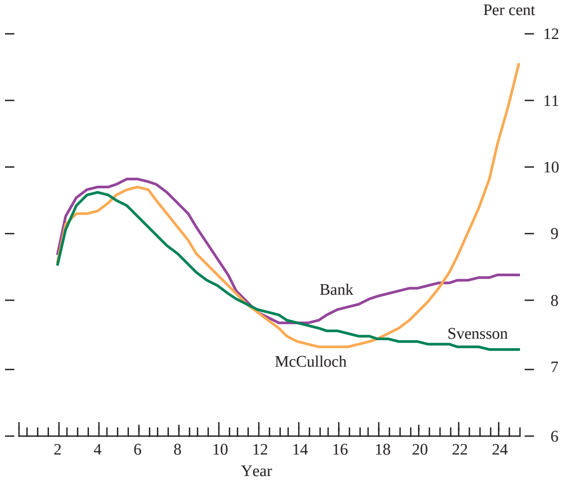

(a) Based on prices on 7 June 1994.

which are not rejected by the data—are desirable, and that the functional form suggested by Svensson is that best suited to provide stability in the shapes of the curves. By imposing these restrictions, of course, the model becomes less able to represent observable data, but statistical tests are unable to distinguish between Svensson's approach and other models. It has therefore been decided that, since the gain from restricting the curves to 'reasonable' shapes is offset by only a small loss in precision, a combination of Svensson's functional specification of the discount function and the current Bank adjustment for tax effects should be the main contender to supplant the Bank's current methodology.

A further decision is required concerning the estimation criterion: should price errors or yield errors be minimised? In principle, smaller yield errors might be preferable, on the grounds that it is the interest rates rather than the prices that are the principal factor for monetary policy purposes; moreover, minimising yield errors implicitly gives greater weight to shorter maturity (and therefore more relevant) bonds. Statistical tests, however, point in the other direction: there appears to be a statistically significant estimation cost to minimising the sum of squared yield errors.

Deriving inflation expectations using index-linked gilt prices

The prices of index-linked gilts, which were first issued following the 1981 Budget and now make up approximately 15% of the UK government bond market, can be used in several ways to derive the markets' expectations of inflation. Index-linked gilts are designed to give the investor a known real return independent of the inflation rate: both the coupon payments and the redemption payment are revalued to keep pace with RPI inflation, so preserving the real value of both income and capital. Index-linked gilts do not offer *complete* real value certainty, however, since there is an eight-month

(1) 'Shorter' in this context means up to about ten years. It seems unlikely that participants form anything more than very approximate expectations about economic variables beyond this horizon.
(2) Since incorporating the Bank's tax model makes the McCulloch and Schaefer approaches to modelling the term structure virtually identical, only the former was included in the testing.

lag in the indexation.[1] The effect of this is that an investor will gain if, once the nominal value of a payment is fixed, inflation falls over the succeeding eight months, and will lose if it rises.[2]

Index-linked gilts allow real rather than nominal returns to be measured and so, in conjunction with the nominal returns estimated from the prices of conventional gilts, allow inferences about inflation expectations to be drawn. But an important consequence of the eight-month indexation lag is that, when computing the real yield on an index-linked gilt, some assumption must be made about future inflation in order to value some of the future cash flows.

Simple measures of inflation expectations

Central to the derivation of inflation expectations from bond prices is the *Fisher identity*. This states that the nominal yield on a bond can be separated into (at least) two components: its real yield and the *average* expected inflation rate. Using a simple interpretation of the Fisher identity, a measure of *average* inflation expectations can be calculated by subtracting the real yield (at some assumed average inflation rate) on an index-linked gilt from the nominal yield on a conventional gilt, preferably of identical maturity. For example, by subtracting the real yield on a five-year index-linked stock from the nominal yield on a five-year conventional, this method gives an estimate of average inflation expected over the next five years.[3]

As mentioned above, the gross redemption yield of a bond is dependent on factors other than its maturity—not least on the size of its coupon—so that matching bonds by maturity and ignoring other factors may produce misleading estimates. Conventional bonds are often compared with one another on the basis of their *duration*—a measure that weights each of a bond's cash flows by the length of time before it is received—to standardise the timing of cash flows. It is sometimes suggested that by analogy it is more appropriate to compare conventional and index-linked gilts of similar duration (rather than maturity). However, this assumes that the factors determining the importance of the timing of cash flows are the same for both types of bond. This is reasonable when comparing two conventionals, since the important factor—the risk of a move in the nominal interest rate—is the same for each bond. But when comparing a conventional with an index-linked bond, the risks are not comparable and it is therefore less clear that matching bonds by duration offers any real advantage over matching by maturity.

Since the real yield on an index-linked bond is dependent on an assumed average rate of inflation, the inflation expectation produced by this method depends to some extent on the original inflation assumption: in effect, an inflation expectation is used to estimate an inflation expectation. Comparing, for example, the real yield on 2% Index-linked 1996 with the nominal yield on 10% Conversion 1996 for a recent date, the latter was 7.4%, while the real yield on the index-linked gilt was 4.5% using a 3% inflation assumption and 4% using a 5% inflation assumption. Using this simple method, the inflation expectations using the 3% and 5% inflation assumptions were 2.9% and 3.4% respectively. The problem can, however, be overcome by using the *break-even inflation rate* methodology: this embodies an iterative procedure that solves for the real yield and the inflation expectation simultaneously, and so does not depend on the original inflation assumption. In the above example, the break-even inflation rate is 2.8%.

Implicit in these computations is the crucial assumption that investors require no risk or liquidity premium for holding either index-linked or conventional gilts; or, if they do, that the premia are identical for the two sorts of asset. The assumption implies that, in an equilibrium where there are no arbitrage opportunities,[4] a conventional and an index-linked stock will have the same expected nominal rate of return.

Problems with simple measures

There are several deficiencies with these methods of deriving inflation expectations. First, it will often only be possible to find pairs of gilts of *approximately* the same maturity, introducing inaccuracies into the values calculated for the real rate and the expected inflation rate. More seriously, there may not be an index-linked stock of even approximately the maturity for which it is wished to derive an inflation expectation.

Another problem is that, since the value of a bond to an investor depends on his or her marginal tax rate, some assumption must be made about the tax rate in order to calculate the return. The tax assumption then feeds through into the calculated inflation expectation. Chart 3 shows the sensitivity to the tax assumption of a break-even inflation rate (calculated here using the 2.5% Index-linked 2003 stock and a conventional stock of similar maturity). It shows how expectations of the average inflation rate over the period until 2003—as measured by the break-even rate methodology for investors facing different marginal tax rates—changed over the course of 1993. Without a view on the appropriate tax rate to apply, it is clear that little useful information can be gained from the *level* of a break-even rate series. It seems, however, that the *changes* in the series vary little with tax: there is a fairly stable differential between the break-even time series at different tax rates.

In addition, the fact that a break-even inflation rate is derived from only two gilt prices—one index-linked and one

(1) When a bond is traded between coupon payment dates, the seller foregoes the next coupon. It is market practice for the buyer to compensate the seller by paying *accrued interest* (over and above the quoted price)—an amount proportional to the period for which the seller has held the bond but will not receive a coupon payment. To compute the accrued interest payment, the size of the next coupon payment must be known. Since coupons on index-linked gilts are paid every six months, and the RPI for a particular month is known only with a lag, a lag of eight months in the indexation of payments is needed to ensure that the nominal value of the next coupon is always known.

(2) The Bank of England's recent publication, 'British Government Securities: The Market in Gilt-Edged Securities', gives more information on index-linked gilts (see in particular Chapter 3). It is available from the Bank of England, PO Box 96, Gloucester, GL1 1YB.

(3) The measure is often misinterpreted as giving an expectation of inflation in five years' time, rather than the average rate over the next five years.

(4) In the context of this article, *arbitrage* involves the simultaneous purchase and sale of two financial instruments for a riskless profit; the equilibrium referred to is a state in which no such opportunities exist.

Chart 3
Break-even inflation rates

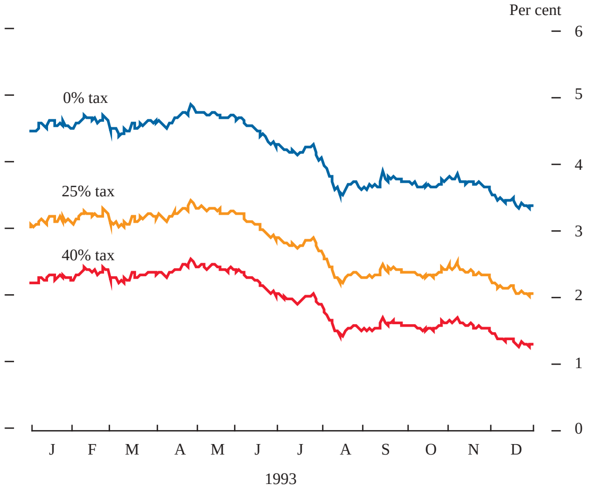

conventional—means that it is particularly vulnerable to distortions produced by the specific pair of stocks selected. For instance, when matching stocks by maturity there may be two conventionals of roughly equal maturity but widely-differing coupons. The difference in the break-even rates derived using the different stocks can be significant, as Chart 4 shows. Another weakness of the approach is that, by

Chart 4
Break-even inflation rates

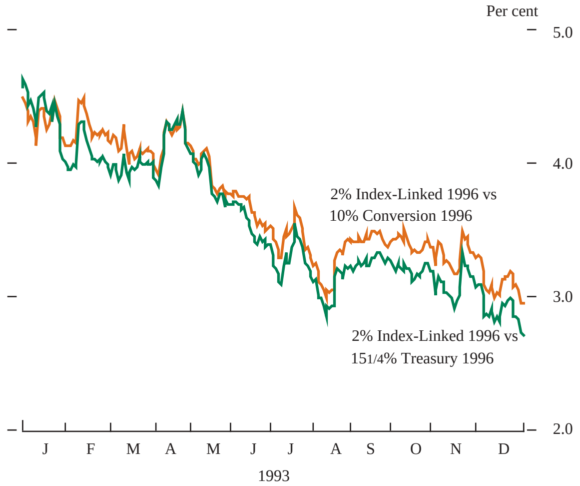

concentrating on only two stocks, it ignores any inflation information contained in the prices of other bonds.

By calculating a break-even inflation rate for each index-linked gilt, it is possible to build up a picture of market expectations of inflation over different time horizons. But as the index-linked market currently consists of only 13 stocks, spread over a maturity range from two to 37 years, it is not very detailed.(1) In addition, such an approach does not allow reliable estimates of implied future one-year inflation rates (as opposed to the average rates that break-even rates represent).

So although these simple measures are useful in showing how inflation expectations may have changed over time, only a limited amount can be learnt from them about the level of inflation expectations.

Term structure of real interest rates

Before looking at the use of term structure models to derive estimates of inflation expectations—an approach which deals with several (though not all) of the problems discussed above—it is helpful to outline how estimates of a real yield curve can be derived using index-linked gilt prices. The estimation of such a curve provides the real equivalent of the nominal interest rate curve discussed above. In particular, it allows a real forward rate curve to be derived. In practice, however, there are two factors which complicate the estimation: there is—once again—the eight-month lag in indexation; and there are far fewer index-linked gilts in issue. The first means that, without some independent measure of expected inflation, real bond yields and hence the term structure derived from real yields are dependent to some degree on the assumed rate of future inflation.(2) The second problem is more practical: there are currently only 13 index-linked bonds in issue and the Bank yield curve model—to take that example—estimates 12 parameters. So using the Bank model as it stands to estimate a real yield curve from index-linked bonds would give an exact fit to the yields observed. Such an approach would be impractical since it would lead to highly unstable forward rate curves.

Despite these problems, the four term structure models investigated can be amended to produce real yield curves dependent upon an assumed rate of inflation. The *Bank model* can be adapted to produce a real yield curve by ignoring all tax effects and simply fitting the yield to maturity structure, once real yields have been calculated for some assumed rate of inflation. In addition, the number of knot points defining the cubic spline can be reduced (from six to three at present) to accommodate the relative lack of data—with little loss of accuracy, since one would not expect the real curve to be as flexible as the nominal. This fitted curve is interpreted as the real par yield curve, from which the term structure of real interest rates and the implied real forward rate curve can be calculated.

The main drawback with this approach is that it includes no parameters to account for taxation effects. Not only is it impractical to apply the full Bank model but it may not even be appropriate, given the difference in nature between the index-linked and conventional markets. The variation of *coupons* on index-linked bonds is not so large as in the conventional market, so tax rules are unlikely to affect to the same extent the prices of indexed bonds with the same maturity but different coupons. But indexed bonds with different *maturities* may attract different categories of investor. Anecdotal evidence suggests that high-rate income tax payers, who are attracted to indexed gilts because of the advantageous ratio of capital to income, prefer the liquidity

(1) In fact, since the index-linked market contains three stocks which are beyond the maturity of the longest conventional, only ten break-even observations are possible.
(2) This dependence can be important when calculating the yields on index-linked gilts approaching maturity, but becomes less important the longer the maturity of the bond.

and reduced price volatility of short-dated securities. In contrast, long-dated index-linked gilts are favoured by pension funds, which are exempt from income tax.[1] The Bank's current implicit assumption is that the marginal investor at all maturities in the index-linked market does not pay income tax, so any distortions introduced by this assumption are likely to be at short maturities. Research continues in this area.

McCulloch's term structure model can be adapted in a reasonably straightforward manner to produce a real yield curve. For the reasons outlined above, however, his tax treatment may not be appropriate for the index-linked market. But estimation of the parameter does at least partially allow for any tax effect that may exist in the index-linked gilt market, so it may still be desirable to include it in the model—even if the estimated parameter cannot readily be interpreted. As with the Bank model, the number of estimating functions is reduced to accommodate the limited number of observations available to be used in the estimation.

Schaefer's model is more difficult to apply, since the lack of data will severely reduce the number of efficient bonds. The data-set therefore needs to be expanded to include all index-linked bonds, in which case Schaefer's approach becomes essentially equivalent to McCulloch's. The *Nelson and Siegel model* has only four parameters, and so can be applied directly to the index-linked market; increasing the flexibility of the model by adding the two extra (Svensson) parameters seems unnecessary.

Implied forward inflation rate curve

Because it is possible to estimate two interest rate term structures for the gilt market—a real yield curve modelling the index-linked sector and a nominal curve modelling conventionals—it is possible to create pairs of hypothetical conventional and index-linked bonds of *identical* maturity[2] for any desired maturity.[3] The break-even approach can then be applied to these pairs to give continuous curves for both average and (more importantly) forward inflation expectations (ie implied future one-year inflation rates). As the prices of most bonds are used in the estimation of the yield curves,[4] this approach has the additional advantage of using virtually all the information on inflation expectations available in the gilt market. It also ensures that the rates derived should be free of any stock-specific distortions and adjusts for most tax effects.

Since an inflation assumption is needed to estimate a real yield curve, the implied forward inflation rate curve[5] that is derived will depend on this assumption—the same problem of consistency that arose with the simple measures of inflation expectations discussed above. To remove this dependency, an iterative procedure has been developed to avoid the need for an assumed inflation rate—in simple terms, the real yield curve is re-estimated for each iteration until consistency between the assumed and the estimated forward inflation rate curve is achieved.

Charts 5 and 6 illustrate implied forward inflation rate curves for the Bank, McCulloch and Svensson approaches. Although there are some differences between the curves, they are all of broadly the same shape; the rise in the McCulloch curve for long maturities is attributable to the lack of constraints that the method places on nominal interest rates.

Chart 5
Implied forward inflation rates[(a)]

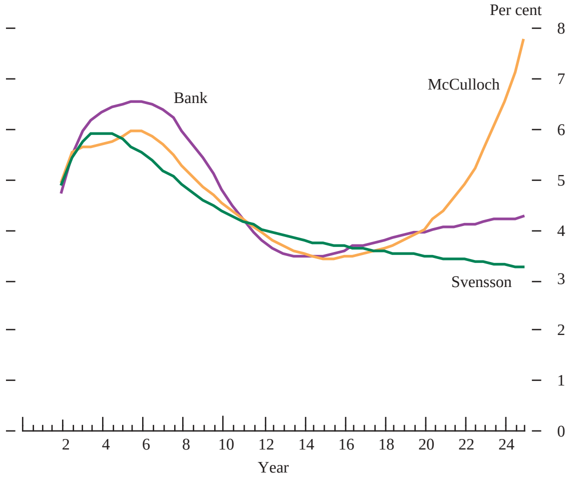

(a) Based on prices on 7 June 1994, so that at the start of year 6, for example, the curves estimate the expected rate of inflation on 7 June 2000.

Chart 6
Implied forward inflation rates[(a)]

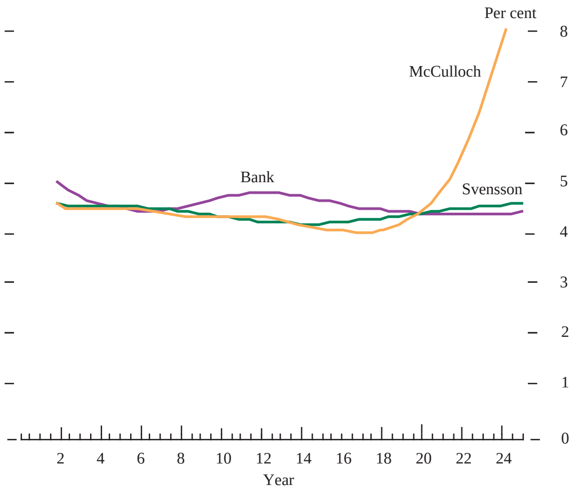

(a) Based on prices on 22 June 1994, so that at the start of year 6, for example, the curves estimate the expected rate of inflation on 22 June 2000.

(1) Pension funds also find long-dated index-linked gilts attractive because they have long-dated, (quasi) index-linked liabilities.
(2) Since the zero-coupon curve is used, the two hypothetical bonds also have identical duration.
(3) It is, however, unwise to extrapolate beyond the maturity of the longest actual bond of either type.
(4) Currently, callable bonds are included in the estimation procedure for nominal curves with a simple rule to determine the maturity date, but convertible, floating-rate, irredeemable and illiquid stocks are excluded.
(5) In past *Inflation Reports* (and in the working papers on which this article is based), the implied forward inflation rate curve is referred to as the inflation term structure.

Since estimation of the forward inflation rate curve relies on a no-arbitrage condition, it is necessary to produce nominal and real curves for the same category of investor. The Bank's current methodology involves modelling both nominal and real curves from the perspective of a zero-rate tax-payer, estimating the nominal curve by modelling the tax effect on the prices of bonds not naturally held by such investors and in this way adjusting them to be comparable with the remainder, and assuming for index-linked gilts that their prices are set by zero-rate tax-payers. The no-arbitrage condition can then be applied, since the index-linked and conventional markets are being compared from the perspective of the same investor paying no income or capital gains tax.(1)

Several previous academic studies have adopted a similar approach to the extraction of inflation expectations. They have not, however, modelled the real yield curve and as a result have derived inflation expectations only for the maturities at which index-linked gilts exist, rather than using the continuous term structure produced by the Bank method. They have also tended to be rather simplistic in dealing with tax effects.

One prerequisite that it may be desirable to impose on a model of the implied forward inflation rate curve is that it is flat for long maturities. This can be achieved by requiring both real and nominal forward curves to be flat at long maturities.(2)

Problems with expectations derived from yield curves

In order to interpret the derived curves as representing 'true' market expectations of interest rates and inflation, it is necessary to assume that the forward interest rates calculated from the term structure model are expected future short rates. In practice, however, there are three main kinds of factor which make it likely that there are differences between the two: those related to risk premia; to liquidity premia; and to 'Jensen's inequality'.

Risk premia

There are two main sources of risk for the holders of government bonds: the risk of unexpected changes in inflation; and the risk of unexpected changes in the spot interest rate. Inflation risk is incurred by holders of bonds with *variable* real returns, ie without a *guaranteed* real return. The inflation risk premium on index-linked bonds is therefore likely to be small, since they offer a high degree of real value certainty. For conventionals, however, it may be significant, because all payments are fixed in nominal terms. The *interest rate* or *price risk premium* represents the compensation a bondholder requires for variability in the value of the bond over time. As the prices of long-duration bonds are generally more sensitive to a change in interest rates than short-duration bonds, the price risk premium included in their returns will be higher.

Liquidity premia

Liquidity premia are important in two respects. First, the prices of bonds that are identical in all respects other than their liquidity may differ. This is particularly likely if one of the bonds is perceived by the markets as a 'benchmark'. How such effects should be treated depends partly on how they are viewed: if, for example, a bond's price is relatively high because the bond is more liquid than comparable stocks, then it is likely to represent the market better. If, however, the bond is being used primarily as a hedge instrument, its price is likely to reflect more than just the term structure of interest rates. The second important effect is a result of the relative liquidity of conventional and index-linked gilts. Because the index-linked market is less liquid, any comparison between the two will implicitly include a premium reflecting the difference in liquidity. The Bank's current methodology assumes that all liquidity effects are negligible, but this may be unrealistic.

Jensen's inequality

There are several competing hypotheses on the economic relationship which should hold between expected future rates and bond prices. If the underlying process corresponds to either of two of the most influential (the *return-to-maturity* hypothesis and the *local-expectations* hypothesis),(3) implied expected future rates will not correspond to actual expected future rates, but will be lower—to an extent dependent on the volatility in future rates. This is essentially because of the difference between $E[(1+r)^{-1}]$ and $(1+E[r])^{-1}$, where $E(.)$ is the expected value function and r is a future interest rate—an example of *Jensen's inequality*.(4)

The effects of both Jensen's inequality and risk premia need to be taken into account when deriving estimates of inflation expectations. The nominal inflation risk premium, the nominal interest rate risk premium and the effect of Jensen's inequality on real rates all tend to bias estimates of inflation expectations upwards. The real inflation risk premium (which is likely to be small), the real interest rate risk premium and the effect of Jensen's inequality on nominal rates work in the other direction.

Preliminary investigation into the effect resulting from Jensen's inequality suggests that it is unlikely to be large if

(1) These estimates can be scaled in the usual way for investors facing other tax treatments.
(2) The restriction is achieved for the Bank model by constraining the cubic spline to flatten at the long end. Although McCulloch's cubic spline will not in general produce asymptotically flat forward rate curves, it can be constrained to do so by applying a technique due to Vasicek and Fong. The McCulloch-based inflation term structures shown in Charts 5 and 6 use the original McCulloch spline specification without the Vasicek and Fong adjustment, and illustrate the unrealistic long rates which this can generate. The functional form of Nelson and Siegel (and that of Svensson) is specifically designed to produce forward rate curves with horizontal asymptotes.
(3) The return-to-maturity hypothesis suggests that the expected return from a single n-period bond is equivalent to the return from rolling over a series of n one-period bonds. The local expectations hypothesis suggests instead that the expected rate of return on any bond in a single period is equal to the corresponding short rate of interest.
(4) More precisely, Jensen's inequality states that for a strictly convex function, the expectation of the function of a random variable will be greater than the function of the expectation of the variable, ie $E[g(x)] > g(E[x])$. For a detailed exposition of the effect of Jensen's inequality when estimating expected interest rates, see the forthcoming *Bank of England Working Paper* by Anderson, N L and Barr, D G, 'Jensen's inequality and the implied forward rate curve'.

the short-term interest rate follows a strongly mean-reverting process. Although it is not clear at this stage how large the *net* effect of risk premia and Jensen's inequality is, it is possible that the Bank's current estimates *overstate* 'true' inflation expectations when referring to implied forward inflation rates in the *Inflation Report* (the risk premia effects probably outweigh those resulting from Jensen's inequality). Further work is under way to model these effects more accurately. Most similar academic studies do not attempt to estimate risk premia.

Summary

Over the past two years, the Bank has investigated a number of sophisticated techniques to extract information on interest rate and inflation expectations from gilt prices. It is not at all easy to choose between them but, on the basis of economic prerequisites and comparative statistical testing, the approach proposed by Svensson seems the most appropriate for creating the forward inflation rate curves used in the *Inflation Report* and for the analysis of monetary policy choices. No final decision has yet been taken on the precise specification to be used to implement the Svensson approach. This—in particular, the choice between minimising price and yield errors—will be determined after further testing and after monitoring the stability of the various options over the coming months. The intention is to make a final decision in time for the new approach to be adopted in the November *Inflation Report*. In order to make as well-informed a decision as possible, the Bank would welcome practitioner comment on the issues raised in this article and on its provisional conclusions.

Irrespective of the precise methodology, it seems clear that further research will be required into the accuracy with which these methods represent market expectations: in particular, the lack of quantification of the effects of risk premia and of Jensen's inequality are potentially important gaps in knowledge. Nevertheless *movements* in the term structures generated using these techniques can be useful indicators of changes in markets' perceptions of policy credibility and so have a role to play in informing monetary policy decisions.