The information content of the inflation term structure

by

*Francis J Breedon***

and

*Jagjit S Chadha**

**	Lehman Brothers, Bishopsgate - this work was carried out whilst on the staff of the Bank of England.
*	Monetary Assessment and Strategy Division, Bank of England, Threadneedle Street, London, EC2R 8AH.

The views expressed in this paper are those of the authors and are not necessarily those of the Bank of England. We would like to thank Mark Deacon, John Lumsden, John O'Sullivan and Paul Wesson for their help and Bruce Devile for invaluable research assistance. The comments of Michael Joyce, Lars Svensson, Paul Tucker, Peter Westaway, Creon Butler and seminar participants at the Bank of England improved the paper. A version of this paper was presented at the 1996 Royal Economic Society Conference, the Federal Reserve Board and Reserve Banks of Atlanta, Chicago and St Louis. Of course any errors and omissions remain ours.

Issued by the Bank of England, London, EC2R 8AH to which requests for individual copies should be addressed: envelopes should be marked for the attention of the Publications Group. (Telephone: 0171-601 4030).

ISSN 1368-5562

Contents

Abstract

This paper examines the inflation forecasting performance of the inflation term structure (ITS) derived by the Bank from conventional and index-linked bonds. We find that the ITS: (i) gives a somewhat better measure of market expectations of inflation than the nominal yield curve alone and (ii) is at least comparable with published inflation forecasts from a number of commercial and academic institutions. But inflation forecasts from the ITS have a tendency to overpredict the level of future inflation which has meant that, in the period we examine, index-linked bonds have proved, *ex post*, to be a cheap form of funding relative to nominal bonds. This overprediction may arise from either an inflation risk premium and/or expectational errors - we cannot be sure which on the evidence.

1 Introduction

Recently, there has been renewed interest in issuing index-linked bonds[1] (see The Economist (1995) and Campbell and Shiller (1996)) and a number of governments are either actively considering or have just started issuing them - New Zealand, Sweden and most recently the United States.

There are a number of advantages which may arise from the issuance of index-linked bonds. Three arguments are summarised here. First, if the private sector perceives that the government has an incentive to create inflation as a means of reducing the real value of outstanding nominal debt, then index-linked funding should reduce the perceived incentive and increase monetary policy credibility (see Calvo and Guidotti (1990)). Secondly, if the government wishes to smooth real expenditure, the interest payments on index-linked bonds are more stable in real terms than conventional debt of the same maturity (see Barro (1995)). Thirdly, if investors charge an inflation risk premium for conventional debt, then the government should be able to achieve lower funding costs by issuing inflation-proof bonds (see Fischer (1975)).[2]

There may also be informational advantages in issuing index-linked bonds. Index-linked bond prices should, theoretically, reflect market participants' expectations of real interest rates and it should therefore be possible to use conventional and index-linked bond prices to derive timely market expectations of inflation. Given this point, it may also be the case that these market expectations provide good forecasts of inflation.

This paper explores some of the evidence for this latter set of benefits which has accrued from the UK index-linked market. Specifically, we consider: whether market expectations of inflation derived from indexed bonds are an improvement on those derived from the conventional nominal curve; how inflation forecasts from the inflation term structure compare with those derived from macroeconomic models; and, whether investors pay a premium for the inflation protection from index-linked bonds.

---

(1) Indexed to some measure of consumer prices.

(2) For example, the primary objective of UK debt management is 'to minimise over the long term the cost of meeting the Government's financing needs, taking account of risk, whilst ensuring that debt management policy is consistent with monetary policy'. See HM Treasury and Bank of England (1995).

The paper is organised as follows: Section 2 describes how data on inflation expectations are derived from the UK index-linked bond market. Section 3 compares the inflation forecasting performance of the inflation term structure with that of the nominal yield curve. Section 4 compares inflation forecasts with those derived from macroeconomic models. Section 5 explores why the inflation term structure tends to overpredict the level of future inflation. And Section 6 concludes.

2 Inflation expectations and index-linked bonds

Index-linked bonds (linked to the retail prices index) were introduced in the United Kingdom in March 1981[3] and were made available to all market participants in 1982 having been sold and traded only by 'eligible' investors - pension funds - before then. By the end of 1982 there were eight in issue but the number soon stabilised to around the current figure of thirteen. It is the market for these bonds that allows the derivation of a real term structure.

Deriving the real term structure

Deriving an inflation term structure from the nominal returns from conventional debt and the real return from index-linked bonds presents two practical problems. First, the small number of index-linked bonds in the market means that it is rarely possible to match the maturity of index-linked bonds and conventional bonds precisely.[4] Second, there is an eight-month indexation lag[5] which means that some neutral assumption about probable inflation is required.[6] In earlier Bank work, Deacon and Derry (1994) overcame this problem by iterating between an initial assumed future inflation rate in order to calculate a real yield curve and the consequent implied

---

(3) Index-linked National Savings Certificates had been in existence since 1975. There have also been a number of private sector issues beginning with the Halifax Building Society in 1985; these private issues are not used in our analysis.
(4) Although reasonable estimates may be derived this way. See Arak and Kreicher (1985).
(5) The combination of publication lags in the retail prices index and the desire to ensure that the next semi-annual coupon is always known with certainty (so that accrued interest can be calculated for tax purposes) means that the bonds are not perfectly indexed and instead are indexed to prices eight months prior to the relevant payment (ie coupon or principal).
(6) Although this eight-month period may seem unimportant, Barr and Pesaran (1994) argue that more of the variance of index-linked bond prices can be ascribed to the effect of changes in expected inflation over this period than to changes in expected real interest rates. Copeland and Levin (1993) even attempt to use this indexation lag in order to estimate the inflation term premium.

inflation forward - the difference between the nominal and real yield at a given fitted maturity - until an inflation profile was found which equalised the inflation assumption and implied inflation expectation.

The estimation of the real yield curve begins by assuming an inflation rate for the first eight months and using this to scale the observed real yields on index-linked debt. These yields are then used to estimate the real yield curve using a simple real forward curve - which is a simplified form of the Nelson and Siegel (1987) methodology - of the form:[7]

$$f(m) = \beta_0 + \beta_1 \, exp(-m/\tau) \qquad \textbf{(1)}$$

where $f(m)$ is the forward rate at maturity, m, and β_0, β_1 and τ are the parameters to be estimated. This highly simplified form of the real forward curve was found to be sufficient to give a close fit to the data since real yields tend to converge to a constant level at relatively short maturities. In effect this means that the real yield curve is flat over all but the shortest maturities. This real forward curve then translates into the following real discount function (ie the relationship between maturity and discount rate) which is the equation that the Bank currently estimates:

$$\delta(m) = exp\,(-m\beta_0\ /100)exp(-\tau\beta_1\,/100)exp((\tau\beta_1\ /100)exp(-m/\tau)) \qquad \textbf{(2)}$$

To deal with the indexation lag, it is necessary to make an assumption about the rate of inflation over the period for which the bond does not give inflation protection. In practice this involved creating a scaling factor that adjusts each cash flow for this lag. So, for example, the present value of cashflow i equals:

$$PV(C_i) = \delta_i \, . \, C_i \, . \, \Lambda_i \qquad \textbf{(3)}$$

where

(7) This simple functional form is necessary given the limited number of index-linked bonds (particularly at short maturities). However, work by Brown and Schaefer (1995) and others implies that this simple functional form should capture most possible shapes of the real term structure. There are of course a number of different functional forms that could be used to fit the real curve see; for example, Gilbert (1995).

$$\Lambda_i = \begin{cases} \left(1+\pi^e_{t,tdi}\right)\left(\frac{RPI_i}{RPI_b}\right) & \text{if } RPI_i \text{ known} \\ \left(1+\pi^e_{t,tl}\right)\left(\frac{RPI_l}{RPI_b}\right) & \text{if } RPI_i \text{ not known} \end{cases} \qquad \textbf{(4)}$$

$\pi^e_{t,tdi}$ = expected inflation between now and cash flow *i*

$\pi^e_{t,tl}$ = expected inflation between now and last published inflation figure

RPI_i = level of price index used to assess cash flow *i*

RPI_b = base level of price index

RPI_l = latest published level of price index

The scaling factor allows for both the inflation effect that comes from the next coupon six months ahead being adjusted by current inflation and for the fact that the latest published inflation number is up to two months out of date. Given assumptions about expected inflation for these two periods, it is then possible to estimate the parameters of the discount function by minimising squared deviations between the fitted yield on a bond (ie the yield constructed by adding together all the cash flows implicit in a coupon bond) and observed yields.

Of course, the estimated discount function will be influenced by the assumptions made about expected inflation over the period of the indexation lag. In order to ensure internal consistency, an iterative procedure is used to generate these assumptions. Starting with some initial values the real curve is estimated. Then a nominal curve is fitted using the 70 or so conventional bonds. The greater amount of data and greater variation of nominal yields means that a more flexible, but related functional form is used (this functional form is the extended Nelson and Siegel approach suggested by Svensson (1995)):

$$f(m) = \beta_0 + \beta_1 \exp\left(m/\tau_0\right) + \beta_2 \exp\left[\left(m/\tau_0\right)\exp\left(-m/\tau_0\right)\right] + \beta_3 \exp\left[\left(m/\tau_1\right)\exp\left(-m/\tau_1\right)\right] \qquad \textbf{(5)}$$

Using the Fisher relationship, we can combine these curves to give an inflation term structure over the eight-month indexation lag period. This new inflation rate is then used to scale the real yield curve again and this iteration continues until the assumed rate equals the actual rate derived (this normally takes about ten iterations). Once the curves have converged, a final real and inflation term structure can be derived. This process is described in more detail in Anderson *et al* (1996).

Chart 1
UK real interest rates - (five-year zero coupon rates)
Per cent

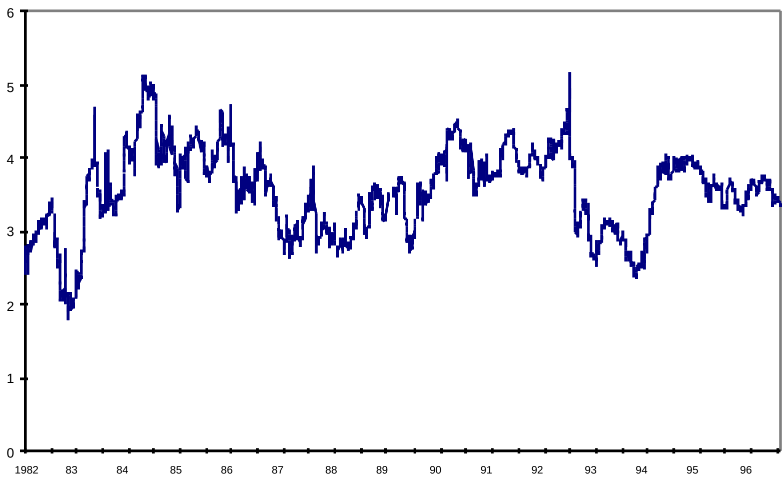

One other adjustment that is made to the nominal curve is an allowance for tax effects. Over the period analysed in this paper many investors were taxed on coupon income but not capital gains. This resulted in the market being characterised by tax clienteles, with higher rate taxpayers favouring low coupon bonds (short selling constraints limited the possibilities for tax arbitrage until January 1996; see Derry and Pradhan (1993)).[8] In the nominal curve this tax effect is allowed for by estimating a set of tax parameters in coupon/capital gain space (see Anderson *et al* (1996) for details). Unfortunately, there are insufficient data to estimate these tax parameters for the real curve so it is assumed that zero-rated institutions are

(8) Short-selling restrictions were removed in January 1996, with the introduction of the gilt repo market.

the marginal investors in indexed gilts. Evidence from the Stock Register tends to support this assumption, as the vast majority of long and medium dated index-linked bonds are held by pension funds, who do not pay tax, and life insurance companies, who do not pay tax on their pension business. However, a significant proportion of shorter-dated index-linked bonds are believed to be held by taxpayers and the extent of the tax distortion this introduces is hard to establish: Robertson and Symons (1993) argue that it is minimal.

3 Measuring market expectations: The inflation term structure (ITS) versus the nominal term structure (NTS)

Using the real and nominal term structures, we can derive an explicit inflation term structure (ITS) - implied market expectations of future inflation. We can now begin to address questions posed at the beginning of this paper. Two points are worth noting at this stage: first, we can only examine the ITS for a relatively short historical period and second the ITS estimate is imprecise because of the effects of risk premia and Jensen's inequality that drive a wedge between bond prices and underlying expectations.(9)

Given that the nominal term structure is a widely used indicator of market expectations of inflation, what value does the inflation term structure add? In a sense, we are simply testing the incremental improvement to yield curve forecasts from using a fitted real curve rather than from making strong assumptions about the behaviour of real yields. Of course all inflation, or for that matter interest rate forecasts from the yield curve contain an expectation of future monetary policy and so may, in general, be inefficient to the extent that there are monetary policy shocks.(10) But all market-based measures of forecasts should, in principle, suffer from the same problem. Our yardstick is simply to compare the forecasting properties of the various measures on the assumption that the one with the better forecasting ability is the better measure of underlying expectations. This is because it seems unlikely that

---

(9) Any attempt to infer market expectations of inflation from the yield curve will need to confront this possibility.

(10) And to the extent the monetary authorities feed back from the yield curve any reduced form estimate will suffer from simultaneity bias; see McCallum (1994).

errors in measuring expectations can do anything but worsen the forecasting performance.[11]

Although there are a wide selection of methods that attempt to derive market expectations of inflation from the NTS[12] we have chosen to focus on just two: the most straightforward - that of Mishkin (1990a, 1990b, 1992) - and one of the most complex - Frankel (1982) and with Lown (1994).

Mishkin's approach is simply to look at the ability of the slope of the nominal term structure to predict future changes in inflation (see equation **8**), implicitly assuming a constant real rate. Frankel, on the other hand, derives an explicit measure of long-run inflation expectations derived from a simple model of the term structure. More precisely, Frankel assumes that the short-term nominal interest rate is expected to move toward a long-run rate - the sum of the economy's long-run real rate and the steady-state inflation rate - at a rate according to the following stochastic differential equation (using Frankel's notation):

$$di_t = - \ (i_t - \pi_o^e - r)dt + \sigma dw, \qquad dw.\pi \sim N\left(0, \sqrt{t}\right) \qquad \textbf{(6)}$$

where i_t is the short term interest rate, π^e is the expected long-run inflation rate and r is the long-run real interest rate which is assumed to be constant; the right-hand term of **6** defines the stochastic process for inflation, which determines the rate of change of the short-term interest rate. Given this process for interest rates and using the expectations theory, Frankel derives the following equation for estimating the expected change in inflation using the slope of the yield curve:

$$i_t^\tau = B0_t + B1_t \left[1 - \left(1 - \exp\left(-\delta\tau\right)\right)/\delta\tau\right] \qquad \textbf{(7)}$$

The measure of the expected change in inflation is the coefficient $B1$ from this regression on bond yields of a range of maturities τ (in years) and a given

---

(11) Of course, even if the inflation forecasts using the real term structure offered very little improvement, they may be preferred because their derivation does not require strong theoretical assumptions about the behaviour of real rates.

(12) See Reichenstein and Elliott (1987) for a range of possibilities. There are many other possibilities, short-term surveys of inflation forecasts could also be used to derive a short-term real structure.

. To derive our measure we use zero coupon yields[13] at semi-annual maturities derived from the yield curve described above and a value of of 0.4.[14]

To compare these measures, we use the simple inflation change equation. By focusing on changes we are able to side-step the problems associated with the analysis of non-stationary variables.

$$\pi^l_t - \pi^s_t = \alpha + \beta [\iota^l_t - \iota^s_t] + \eta_t \qquad \textbf{(8)}$$

π_t = Average inflation between period l (or s) and period t (where $l > s$)
ι_t = Zero coupon rate at period t of maturity l (or s).[15]

This equation allows us to assess (i) whether the predictions from various measures of expectations have significant predictive power, where β is significantly different from zero and (ii) whether they are consistent, where $\beta = 1$. If the null hypothesis that $\beta = 1$ is rejected, then this implies, in the presence of rationality, either incorrectly modelled real rates and/or that time-varying term premia exist which are correlated with inflation outturns. One small extension that we make to this testing framework is to include current inflation as one of the π^s's used. This has the advantage that we can test how well our measure predicts the future changes in inflation from its current level rather than its unknown level some time in the future.[16] Note that in these cases we also substitute current inflation for the ι^s in the NTS and ITS cases: these cases are denoted as $s = 0$.

Before estimation there are several econometric issues to consider. First, as is well known, the overlapping forecast horizons generated in this test give rise to a moving average error process in the residuals of equation **8**. To allow for this we estimate robust standard errors using the Newey-West procedure.[17]

---

(13) Frankel uses yields on coupon bonds but the theory is expressed in terms of zero coupon rates.
(14) This value was found through a maximum likelihood search procedure.
(15) Note that $\iota^l - \iota^s$ is replaced by the single estimated slope variable in the Frankel approach.
(16) The practical advantage for policy of this measure is clear. An indicator which can tell you what inflation will be in two years time conditional on knowing what it will be in one years time is less valuable than one that can tell you what inflation will be in one or two years time relative to its current level.
(17) We also experimented with using IV estimation to overcome potential measurement error problems, see Mankiw (1986), but this made little different to the results.

Also, as our empirical estimates are constrained to the period over which an index-linked bond market developed in the United Kingdom (since 1982), the β estimates are subject to the possibility of small sample bias. To evaluate the extent of any bias, Table A also reports results that are significantly different from zero at the 5% level in Monte Carlo simulations.

Results

Table A shows β estimates from using the NTS, the ITS and the Frankel measure (FE) from the same set of regressions. The differences in forecasting performance are not, in general, statistically significant and this may not be particularly surprising given the relatively short sample period. However, the results do suggest that there are some incremental improvements to inflation forecasting arising from the ITS.[18] The ITS provides significant information (using asymptotic t-values) on future inflation for nine out of the 15 horizons examined, compared with five for the NTS and four for the FE. The Monte Carlo simulations give slightly higher critical values and suggest an equal number of significant parameter estimates for both the NTS and ITS (five), where FE has only three. But the ITS estimates are better determined with lower standard errors, particularly over, arguably, the most important horizons where $s = 0$. Averaging over all 15 horizons, we also find that the ITS is able to explain 32% of the variation in inflation outturns, compared with 23% and 13% for the NTS and FE, respectively.[19]

In general, the performance of the ITS is creditable but not substantially better than the NTS. Are the expectations consistent? As the point estimate of β is greater than one in six of the nine significant cases with asymptotic t-ratios and four of the five cases with Monte Carlo t-ratios this may suggest a lack of consistency. As suggested earlier, this may arise from either an inadequately modelled real rate process and/or the presence of a time-varying risk premium correlated with the set of inflation expectations.

We can assess directly the quality of the modelled real rate process by running equation **8** for the average real rate outturn using the real rate forecasts from the real term structure.[20] The third column of Table B shows

---

(18) A similar result, using a different method, was found by Barr and Campbell (1995).

(19) Recursive estimates of the constant, α in equation **8** , suggest parameter constancy for the ITS equations. Results available from the authors on request.

(20) The dependent variable is created by taking the base rate and, assuming perfect foresight, subtracting the average inflation rate over the next three months.

that we find for the five horizons where $s = 0$ that the estimated slope coefficient is significantly different from zero in all cases and we cannot reject the hypothesis that $\beta = 1$ for three of the five horizons. It would appear then that any lack of consistency may be more likely to result from the existence of a time-varying risk premia or expectational errors.

Table A
β estimates

***l - s* years**	**NTS**	**ITS**	**FE**
1-0	0.81 (2.77)*I	0.64 (3.27)*I	0.54 (5.18)*I
2-0	0.99 (3.64)*I	1.25 (9.10)*I	0.54 (3.81)*I
3-0	1.20 (5.05)*I	1.45 (7.57)*I	0.49 (3.27)*I
4-0	1.35 (6.13)*I	1.45 (6.30)*I	0.41 (2.98)*
5-0	1.38 (6.59)*I	1.34 (5.61)* I	0.44 (1.92)
2-1	0.53 (0.92)	-0.17 (-0.52)	0.01 (0.08)
3-1	0.99 (1.36)	-0.01 (-0.03)	-0.05 (-0.58)
4-1	1.07 (1.32)	0.22 (0.69)	-0.12 (-1.63)
5-1	1.10 (1.75)	0.44 (1.66)	-0.16 (-1.07)
3-2	0.95 (1.35)	0.21 (0.41)	-0.06 (-1.07)
4-2	0.96 (1.13)	0.60 (1.25)	-0.13 (-1.55)
5-2	0.96 (1.48)	0.89 (2.38)*	-0.20 (-1.63)
4-3	0.48 (0.67)	0.95 (2.02)*	-0.07 (-1.79)
5-3	0.55 (0.87)	1.20 (2.37)*	-0.16 (-1.90)
5-4	0.57 (1.16)	1.41 (2.98)*	-0.07 (-1.51)

Notes:
1) Estimation using end-month data 82:4 to 96:9.
2) Newey-West corrected t-ratios in brackets.
3) First column denotes length of period *l* and period *s* respectively used in equation **8** .
4) * = estimates significantly greater than 0 at the 5% level
5) I = estimates significantly greater than 0 at the 5% level using Monte Carlo standard errors. The Monte Carlo simulations are based on ARMA models of the variables estimated from the relevant sample. Errors were i.i.d. and test statistics were produced from 1,000 draws.

Because the point estimate of β is also found to be greater than one when it is found to be greater than zero, it seems possible that there is a time-varying risk premium. But coming directly to such a conclusion is not straightforward. This is because an estimate greater than one implies that any time-varying risk (probably, inflation) premium is *negatively* correlated with the inflation term structure and has a higher standard deviation than the ITS.[21]

(21)   See Mishkin (1990a) for a demonstration of this proposition which follows from solving and calibrating the probability limit of β from equation **8** for an omitted variable.

Table B
Estimated constant and standard error with β constrained to one

	Inflation predictions of:		Real rate predictions of:
Horizon: ***l*- *s***	**NTS**	**ITS**	**real term structure**
1-0	-4.87 (0.48)	-1.84 (0.43)	0.80 (0.05)
2-0	-4.95 (0.65)	-1.66 (0.45)	0.95 (0.02)
3-0	-5.19 (0.71)	-1.74 (0.56)	1.02 (0.02)
4-0	-5.48 (0.74)	-1.89 (0.66)	0.96 (0.03)
5-0	-5.47 (0.61)	-1.89 (0.61)	0.95 (0.02)

Notes: Newey-West adjusted standard errors of the regressions in brackets. In the first two columns the estimated constant gives the average real interest rate in the NTS case and the average prediction error in the ITS case and base rate case. The third column reports the β from a real rate forecasting equation.

Such a negative correlation may be considered intuitively unlikely and certainly, most studies of the inflation risk premium in indexed debt have found that it is positively related to inflation. For example, Chu, Lee and Pittman (1995) find that the UK inflation risk premium is positively related to the level of inflation over the period 1985-91. On the second point, it is also difficult to be sure how the relative variances of expected inflation and risk premium match up but Barr and Pesaran (1995) find that expected inflation has a higher variation than their identified risk premia. The seemingly biased estimate may then simply arise from small sample problems and in fact with the Monte Carlo standard errors only in one case is β found to be significantly greater than one. We shall return to the relative importance of expectational errors and risk premium in Section 5, when we discuss the reasons for the overprediction of the level of inflation.

Although the tests used in Table A are informative, the usefulness of the point estimates of β is open to question. This is because if we use the point β estimate, we imply that the observed term structure should be multiplied by the estimated coefficient to give the best estimate of future inflation. Such an adjustment factor would be difficult to justify unless one had a good reason. For practical purposes therefore, the ITS data would tend be interpreted assuming that β is equal to one. So we analyse the forecasting properties of the ITS and NTS with this restriction imposed. This allows us to also focus on the cases where current inflation is used as the *s*-period rate - that is the expected change in inflation from its current level - not only because this is the most important benchmark for policy applications but because this constraint also allows us to estimate any significant tendency for our measure to over or underpredict the level of future inflation.(22)

The first two columns of Table B assess the quality of the predictions from the NTS and ITS with β constrained to equal one. The average standard error is substantially smaller in the case of the ITS at 0.54 compared with 0.64 for the NTS. Over horizons other than $s = 0$ the relative performance of the NTS is slightly better (results not shown). In some sense, this result is not surprising because the intercept term in the NTS regression is controlling for two factors, average inflation overprediction and average real interest rates, at 3.4%. The intercept from the ITS equation is simply a measure of the level of inflation overprediction, 1.8%.(23) That the standard error does not fall further provides evidence to suggest that the time series of fitted real rates is relatively stable.

To sum up, the addition of a real term structure seems in general to marginally improve yield curve based inflation forecasts and to improve significantly their efficiency. That inflation forecasts are not a statistically significant improvement may be the result of the relatively short sample of data and the fact that real rates seem relatively less variable than inflation forecasts and inflation.

4 The inflation term structure and forecasts of inflation

Charts 2 to 4 below show that the ITS predictions of future changes in inflation do seem to predict a number of turning points and its performance appears particularly good at longer horizons. But how well the inflation term structure forecasts inflation compared to forecasts from other sources? We might expect that a forecast derived from financial markets should theoretically contain all information available at the time and so should be at least as good as forecasts from any other source.(24) Against this however the

---

(22) This analysis concentrates on the comparison between the NTS and the ITS, as the estimate of β in the case of FE is already constrained by assumption to lie between zero and one.
(23) The fact that the intercept in the ITS regresion only measures inflation overprediction but that the standard error only falls by around 0.10 suggests that real rates are not especially variable.
(24) Except possibly agencies that have access to private information relevant to the inflation process.

Chart 2: One-year ahead change in inflation: predicted and actual

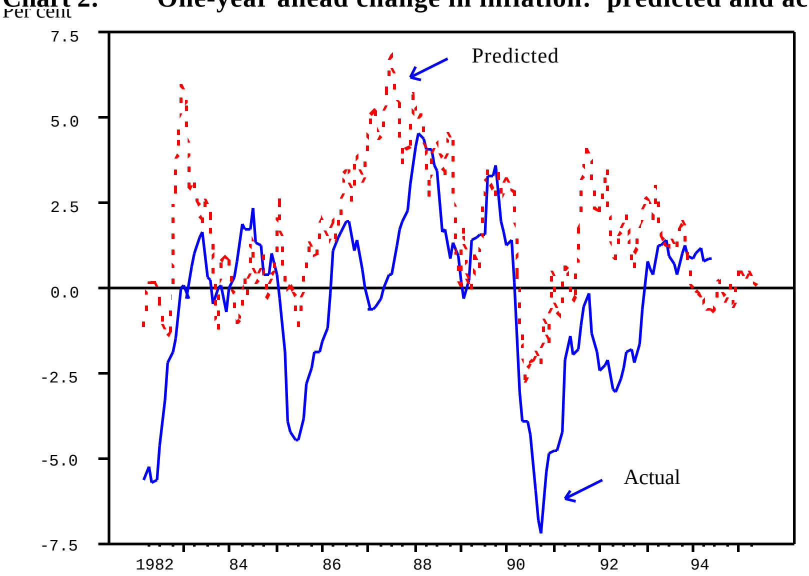

Chart 3: Three-year ahead change in inflation: predicted and actual

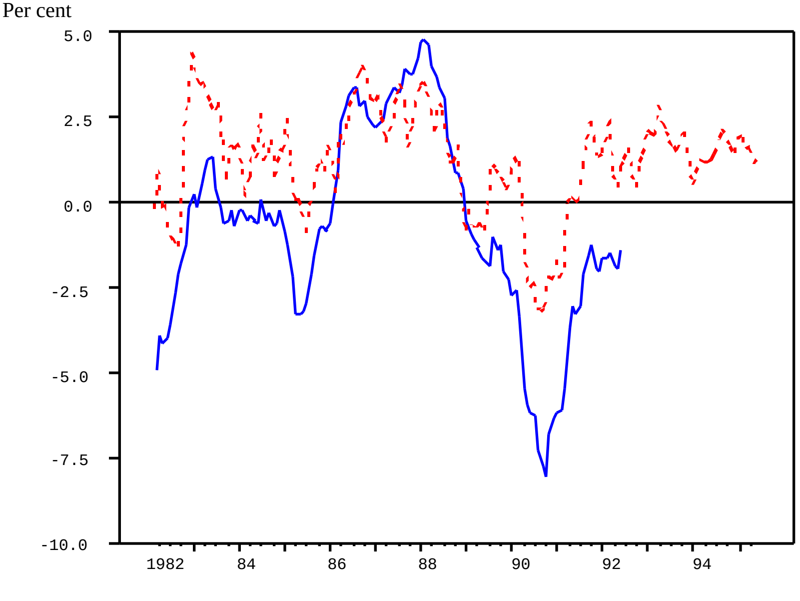

Chart 4: Five-year ahead change in inflation: predicted and actual

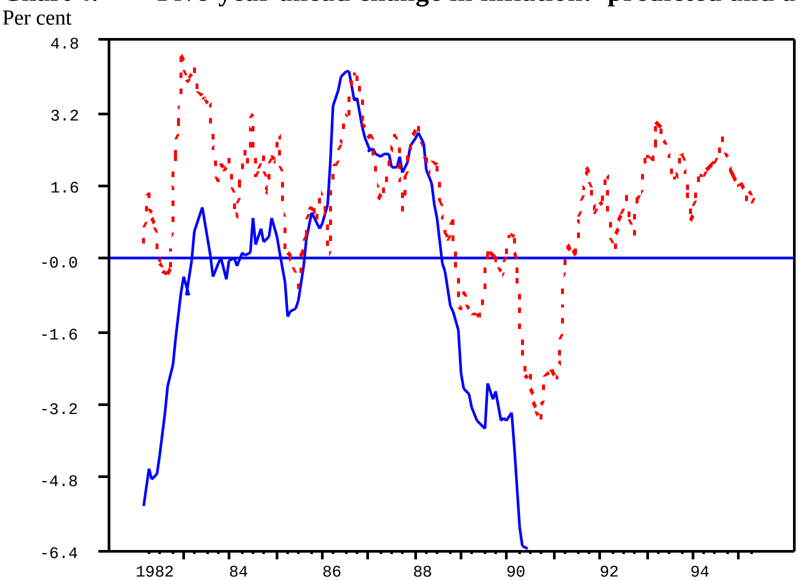

Notes: Charts 2 to 4 show the actual and predicted difference between current annual inflation and average inflation over the forthcoming one, three and five years. The prediction is the relevant zero coupon inflation rate minus current inflation.

earlier results suggest that there may be some bias in the inflation term structure forecasts arising possibly from significant risk premia.

This means that the identified inflation term structure is an imprecise measure of the underlying expectations and so other forecasts may well be superior in practice. And when comparing forecasting performance of the ITS with non-yield curve forecasts differences may even arise because different forecasts may be sampling the expectations of different groups of agents.

In order to assess the relative forecasting abilities of the inflation term structure, we compared them with a database of published forecasts. Our source was Burrell and Hall (1994) who constructed a detailed database of forecasts obtained from *Economic Forecasts- A monthly worldwide survey* back to 1984.[25] The database gives monthly forecasts of inflation at the end

(25) The forecasts from the December issue of this publication are used as the basis of the 'golden guru' award published in the Guardian/Independent on Sunday.

of the following year by a number of commercial and academic institutions. Since the inflation term structure give forecasts on a semi-annual basis, we used only two months out of the twelve available for each year, June (an eighteen-month forecast) and December (a twelve-month forecast). For June we compared the one-year forward rate six months ahead with the published forecasts whilst for December the one-year spot rate was used. This gave us a database of eighteen inflation forecasts; nine forecasts made in June of next year's inflation and nine made in December finishing in December 1992.

Table C
A comparison of the inflation term structure with published forecasts

	Inflation term structure	Average of published	Average of City	Average of academic	HM Treasury
RMSE	2.34	1.95	1.80	2.23	2.17
Mean error (actual - forecast)	-1.71	0.68	0.46	0.93	0.99
RMSE - inflation changes	1.65	1.83	1.74	2.03	1.94

Looking at root mean squared errors (RMSE) in Table C, the inflation term structure gives the worst performance in terms of forecasting ability, but, as the next rows show, this is because, as was already shown in Table B, the systematic bias in its forecast is far larger than that of the published forecasts: this, again, may be evidence of an inflation risk premium. Adjusting the RMSE for this systematic bias (ie focusing on the ability to forecast changes in inflation rather than the level) shows that the inflation term structure marginally outperforms the others shown here - but not all forecasts in the sample.(26)

Although the ITS performs relatively well in the comparison above, there are a number of reasons to expect its longer-term relative forecasting performance to be better than its short-term performance. This is largely due to the problems of the indexation lag which means that the short-term forecasts of under eighteen months will - by construction - to a great extent be an extrapolation of the longer-term one. Added to this, problems with tax effects and liquidity may well influence short-term forecasts. To look at comparative performance over longer horizons we compared the ITS with forecasts by the

(26) The inflation term structure forecast also exhibited higher variance than all other forecasts in the sample though still less than 40% of that of actual inflation.

London Business School's Centre for Economic Forecasting between June 1982 and June 1991.(27) In this case we compared forecasts of the average level of inflation over the next one to four years rather than forecasts of inflation for each year individually, ie zero coupon rates rather than forward rates.

Table D
A comparison of the inflation term structure and London Business School forecasts

	1 year ahead	2 years ahead	3 years ahead	4 years ahead
ITS - RMSE	2.90	2.31	2.08	1.96
Mean error	-2.02	-1.71	-1.57	-1.63
RMSE inflation changes	2.07	1.55	1.36	1.10
LBS- RMSE	1.37	1.59	1.76	1.62
Mean error	0.59	0.80	0.87	0.81
RMSE inflation changes	1.24	1.37	1.52	1.40

Notes: 23 forecasts in the sample, except four year ahead based on a sample of ten forecasts. Forecasts horizons rounded to the nearest quarter.

The performance of the inflation term over shorter-term forecasts is worse over this sample than in the previous sample but, as expected, its relative performance seems to improve with the forecast horizon. Overall it seems that if we focus on predicting changes in inflation rather than its level, the ITS is at least as good as published forecasts, particularly over longer horizons. We also note that its poor performance at predicting the level of inflation is due to a tendency to overpredict that is far larger than the tendency of published forecasts to underpredict inflation over the period analysed. Changes in the ITS may therefore provide a very useful off-model guide to changes in inflation expectations. But we should point out that the comparison is necessarily limited, in terms of both horizon and number of forecasts analysed.

5 Why has the ITS overpredicted the level of inflation?

As Table B showed, although the ITS seems to have predicted *changes* in inflation quite well, it has significantly overestimated the *level* of inflation over our sample period. Sections 3 and 4 also suggest that the extent of this overprediction is seemingly robust to choice of horizon and to sample. In

(27) The only forecasting organisation that has consistently published long-term forecasts in its *Economic Outlook*.

other words, the estimated constant from the ITS regression - which is a measure of the average forecast error - is quite large averaging about 1.70-1.80%. This overprediction translates into higher returns on nominal debt than indexed debt.

There are two possible explanations for the overprediction which are non-exclusive. First, there is the possibility of two related types of expectational error in market forecasts. The inflation term structure may be a good measure of financial market inflation expectations but a poor forecast of actual inflation if market expectations are a biased forecast. This may apply especially, for example, in a period of low credibility when the monetary authorities plan to pursue a policy which differs from what the market expects. If this policy is carried through realised inflation will deviate from expected and *ex post* market forecasts would be biased. The second type of expectational error results from our relatively short sample of market expectations; we may have a sample which does not adequately capture the underlying probability distribution of market expectations and is therefore biased.[(28)] As a result of either or both of these points nominal bonds were too cheap relative to indexed bonds *ex post*.

The second possible explanation is that there is an inflation risk premium. This implies that investors may have knowingly paid more for index-linked bonds since they valued the inflation protection. As Fischer (1975) shows, in the absence of other hedges against unanticipated inflation, this willingness would lead to relatively low *ex post* yields on index-linked bonds. Although our relatively short sample period means that it is not possible the definitively distinguish between these two explanations, there are three pieces of evidence that suggest that the inflation risk premium may account for at least a part of the overprediction.

1) Columns one and two of Table E show the ability of the nominal term structure to predict future base rates[(29)] and the real term structure to predict real rates; both columns give the constant in a regression where β has been constrained to one. We find no significant overprediction of nominal interest rates and that the real curve underpredicts the level of future real rates. The corresponding overprediction of inflation, over this period, seems more related then to the real curve than the nominal curve *per se*. This is the prediction of Fischer (1975) who would ascribe the overprediction to an inflation risk

---

(28) Otherwise known as the 'peso problem'.

(29) The commercial bank rate closely related to the official base rate.

premium. But again, we cannot rule out the possibility that real rates could be underpredicted for expectational reasons.(30)

Table E
Is the level of overprediction the result of the nominal or real curve?

	Nominal rate prediction	Real rate prediction
Horizon: *l-s*	**NTS**	**RTS**
1-0	0.48 (0.21)	2.47 (0.14)
2-0	0.42 (0.46)	2.26 (0.09)
3-0	0.35 (0.66)	2.21 (0.08)
4-0	0.14 (0.77)	1.69 (0.08)
5-0	0.11 (0.68)	1.72 (0.08)

Note: Newey-West standard errors in brackets.

2) Table C also provides some evidence on the nature of the bias in the inflation term structure. Published forecasts tend to underestimate rather than overpredict the level of inflation(31) over this period. If expectations were consistent across sectors this would give some support to the hypothesis that the overprediction by the inflation term structure is related to risk premia and not expectational errors (this interpretation is similar to that of Froot's (1989) analysis of nominal term structure predictions in the United States). But, of course, we cannot rule out the possibility that structural models may be encapsulating the expectations of very different agents from those in financial markets.

3) Although the size of overprediction seems very large for an inflation risk premium, there is some international evidence to support figures in this range. Some estimates have been derived in the United States from analysis of the CPI future that traded on the New York Coffee, Sugar, and Cocoa Exchange between June 1985 and February 1987: Chu (1991) estimated an inflation risk premium on that market of 1.4% whilst Fleskaer and Ronn (1988) estimated the premium to be 0.9% at a one-year horizon. A higher measure risk premium for the United Kingdom may simply result from a more volatile inflation performance than the United States. Certainly, for example in the

(30) That is a correct estimate of the nominal interest rate but consistent errors on the split between inflation and real interest rates.

(31) Interestingly, the degree of underprediction is similar in magnitude to the underprediction of official interest rates by the nominal term structure.

case of Israeli index-linked bonds, Kandel, Ofer and Sarig (1996) find an inflation risk premium of about 0.34% per month over the period 1984 to 1992.

6 Conclusion

This paper has examined the information content of inflation forecasts that can be derived using index-linked bonds and conventional bonds. We find that the derived ITS gives a somewhat better indication of the bond market's inflation expectations than can be derived using either the nominal term structure alone or a variant employing assumptions about real interest rate behaviour. The inflation forecasts of the ITS also seem at least as good at forecasting future changes in inflation as forecasts derived from macroeconometric models. These characteristics of the ITS and its timeliness tend to make its inflation forecasts a useful addition to policy analysis. Because of the real term structure's tendency to underpredict future real interest rates, index-linked bonds have proved, *ex post*, to be cheap funding for the UK government.[32] But we cannot be sure whether this underprediction results from an inflation risk premium or expectational error and also cannot know whether this overprediction will persist.

(32) Particularly, given the fact that the risk characteristics of government index-linked bond issuance appear to be good, see Barro (1995).

References

Anderson, N, Breedon, F, Deacon, M, Derry, A, and Murphy, G (1996), *Estimating and Interpreting the Yield Curve,* London: John Wiley.

Arak, M and Kreicher, L (1985), 'The real rate of interest: inferences from the new United Kingdom indexed gilts', *International Economic Review,* 26, pages 399-408.

Barr, D and Pesaran, B (1995), 'An Assessment of the Relative Importance of Real Interest Rates, Inflation, and Term Premium in determining the Prices of Real and Nominal United Kingdom Bonds', *Bank of England Working Paper*, No 32.

Barr, D and Campbell, J (1995), 'Inflation, Real Interest Rates, and the Bond Market: A Study of United Kingdom Nominal and Index-Linked Government Bond Prices', *Harvard Institute of Economic Research Discussion Paper*, No 1732.

Barro, R (1995), 'Optimal Debt Management', *NBER Discussion Paper*, No 5,327.

Brown, R and Schaefer, S (1995), 'The term structure of real interest rates and the Cox, Ingersoll Ross model', *Journal of Financial Economics*, 35, pages 3-42.

Burrell, A and Hall, S (1994), 'A comparison of macroeconomic forecasts', *London Business School Centre for Economic Forecasting Discussion Paper,* 20-94.

Calvo, G and Guidotti, P (1990), 'Indexation and maturity of government bonds: an explanatory model', in Dornbusch, R, and Draghi, M (eds), *Public Debt Management: Theory and History,* CEPR, pages 52-81.

Campbell, J and Shiller, R (1991), 'Yield Spreads and Interest Rate Movements: A Bird's Eye View', *Review of Economic Studies,* 58, pages 495-514.

Campbell, J and Shiller, R (1996), 'A Scorecard for Indexed Government Debt', *NBER Macroeconomics Annual.*

Chu, C (1991), 'Futures Prices and Inflation Information', *Review of Quantitative Finance and Accounting,* 1, pages 281-91.

Chu, Q, Lee, F and Pittman, D (1995), 'On the inflation risk premium' *Journal of Business Finance and Accounting,* 22, pages 881-92.

Copeland, L and Levin, E (1993), 'Reading the message from the United Kingdom Indexed Bond Market: Real Interest rates, Expected Inflation and the Risk Premium', *The Manchester School,* 61, pages 13-34.

Deacon, M and Derry, A (1994), 'Deriving Estimates of Inflation Expectations from the Prices of United Kingdom Government Bonds', *Bank of England Working Paper,* No 23.

Derry, A and Pradhan, M (1993), 'Tax specific term structures of interest rates in the United Kingdom government bond market', *Bank of England Working Paper,* No 11.

The Economist (1995), 'Index-Linked Bonds', May 27th.

Fischer, S (1975) 'The Demand for Index Bonds', *Journal of Political Economy*, 83, pages 509-534.

Flesaker, B and Ronn, E (1988), 'Inflation Futures and a Riskless Real Rate', *Review of Futures Markets,* 8, pages 37-76.

Frankel, J (1982), 'A Technique for Extracting a Measure of Expected Inflation from the Interest Rate Term Structure', *Review of Economic Studies,* 61, pages 135-42.

Frankel, J and Lown, C (1994), 'An Indicator of Future Inflation Extracted from the Steepness of the Interest Rate Yield Curve Along Its Entire Length', *Quarterly Journal of Economics,* 109, pages 517-30.

Froot, R (1989), 'New hope for the expectations hypothesis of the term structure', *Journal of Finance,* 44, pages 283-305.

Gilbert, C (1995), 'Alternative models for estimating real interest rates and expected inflation from United Kingdom Index-Linked Bonds', *Queen Mary and Westfield College, London mimeo*.

H M Treasury and Bank of England (1995), *Report of the Debt Management Review*.

Kandel, S, Ofer, A and Sarig, O (1996), 'Real Interest Rates and Inflation: An Ex-ante Empirical Analysis', *Journal of Finance,* 51, pages 205-225.

London Business School, *Economic Outlook*, Gower, various issues.

McCallum, B (1994) 'Monetary Policy and the Term Structure of Interest Rates', *NBER Working Paper*, no 4,938.

Mankiw, N (1986), 'The Term Structure of Interest Rates Revisited', *Brookings Papers,* 1, pages 61-96.

Mishkin, F (1990a), 'What Does the Term Structure Tell us About Future Inflation?', *Journal of Monetary Economics,* 25, pages 77-95.

Mishkin, F (1990b), 'The Information in the Longer Maturity Term Structure About Future Inflation', *Quarterly Journal of Economics*, 105, pages 815-28.

Mishkin, F (1992), 'Is the Fisher Effect for Real?', *Journal of Monetary Economics*, 30, pages 195-215.

Nelson, C and Siegel, A (1987), 'Parsimonious modelling of yield curves', *Journal of Business,* 60, pages 473-89.

Reichenstein, W and Elliott, J W (1987), 'A Comparison of Models of Long-Term Inflation Expectations', *Journal of Monetary Economics***,** 19**,** pages 405-25.

Robertson, D and Symons, J (1993), 'Real interest rates and Index-Linked Gilts', *LSE Centre for Economic Performance Discussion Paper*, No 181.

Rutterford, J (1983), 'Index-linked gilts', *National Westminster Bank Quarterly Review*, November, pages 2-17.

Söderlind, P (1995), 'Forward Interest Rates as Indicators of Inflation Expectations', *Institute For International Economic Studies,* Seminar Paper Number 594.

Svensson, L (1995), ‘Estimating Forward Interest Rates with the Extended Nelson and Siegel Method’, *Sveriges Riksbank Quarterly Review*, 1995:3, pages 13-26.

Yariv, D (1995), ‘Market based inflationary expectations as an indicator for monetary policy: The case of Israel’, *Bank of Israel mimeo*.