EXTRACTING INFORMATION FROM ASSET PRICES IN AN UNCERTAIN WORLD

by

Peter Westaway,

Monetary Instruments & Markets Division,

Monetary Analysis,

Bank of England,

Threadneedle Street, London EC2R 8AH

Phone: 00 44 (0)20 7601 5686

Fax: 00 44 (0)20 7601 5953

Email: peter.westaway@bankofengland.co.uk

Paper prepared for the Second Monetary Policy Research Workshop in Latin America and the Caribbean on Monetary Policy, Uncertainty, and the Business Cycle, a conference organised by the Central Bank of Peru and the Centre for Central Banking Studies (CCBS) of the Bank of England, held in Lima, Peru, November 6-7, 2006.

Abstract:

Extracting information from asset prices in an uncertain world

This paper is concerned with how policymakers respond to the presence of uncertainty. Typically, questions addressing this problem are concerned with how policy should optimally be set in the presence of uncertainty about model parameters or about key aspects of the behaviour of economic agents. This paper is concerned with a different aspect to that question, namely how economic agents themselves respond to the presence of uncertainty. This matters because the policymaking problem is usually examined in the context of economic models where uncertainty is either ignored or subsumed in the constants of the model. This may not lead to serious policy errors if the volatility of the shocks hitting the economy is not changing or if agents' attitude to uncertainty is constant. But if uncertainty is time-varying then our ability to interpret agents' behaviour and the shocks hitting the economy will potentially be compromised if we ignore that feature of the economic environment.

Asset prices represent a rich and timely source of information about current and future economic fundamentals. But they also provide important information about market participants' perceptions of expected volatility in the environment and their attitudes towards it. This is revealed in observed and expected risk premia, the excess return investors earn on risky assets, and also in option prices which allow us to extract information about the expected distribution of asset returns. This paper explains how this information is extracted from a range of financial markets and used to provide policy advice to the Monetary Policy Committee of the Bank of England. It aims to show how we apply our modelling toolkit and how we are continuing to refine the models we use to improve our modelling of risk premia. And it explains how we are following a research agenda which is attempting to incorporate risk and uncertainty into a DSGE modelling framework with the aim of strengthening our theoretical understanding of the links between macro and finance, and hence of how uncertainty impinges on economic behaviour.

Since the Bank of England's forecast is constructed within the framework of a forward-looking general equilibrium model, the paper begins by explaining how the Bank's forecast is conditioned on a market-implied path for short term interest rates and examines the special challenges that asset prices present in a forward-looking forecasting framework.

The paper goes on to describe the range of financial market information regularly presented to policymakers at the Bank of England.

For interest rates, it describes the range of nominal and real yield curves that are estimated daily for the UK, euro area and US yield curves and explains how these are used to derive measures of near-term policy rate expectations as well as measures of real rate and inflation expectations at a range of horizons. Recent research is described which explains how these measures can be decomposed into rate expectations and term premia, so casting light on the underlying causes of the recent bond yield conundrum.

For equities, it explains how the dividend discount model can be used to evaluate the extent to which market valuations are warranted by current earnings expectations and interest rates.

For exchange rates, it explains how news in relative yield curve movements can be used, under the maintained hypothesis of the uncovered interest parity condition (UIP), to account for exchange rate movements in terms of cyclical news, risk premia or movements in the equilibrium exchange rate.

And finally for derivatives, it shows how option-implied PDFs can be derived for a range of financial instruments. And preliminary work is described which is intended to transform these risk neutral PDFs into subjective probability distributions which would give us a more direct reading of market participants' probability distribution of outcomes.

The risk premium plays an important role in all of these modelling tools but often it represents the residual component which reconciles asset price changes with movements in observed fundamentals. It is important to understand whether the apparent movements in risk premia are consistent with our theoretical understanding of how risk premia might reasonably evolve. To examine this question, we have been conducting a research program which uses a DSGE model of (at this stage) a closed production economy to examine how risk premia are determined by the interaction of the shocks impinging on the economy, the preferences of agents and the structural characteristics of the economy itself. So far, this has been used to cast doubt on the argument that the "Great Stability" can explain the fall in bond yields. And it has helped us rationalise the stylised facts about the slope of the real and nominal yield curves.

This theoretical macro-finance agenda has also forced us to think harder about those circumstances when the textbook model of asset price determination breaks down. Such departures from the textbook paradigm are commonplace in terms of market anecdote relating to the effect on bond yields of asset-liability management by pension funds or the build-up of dollar reserves by Asian central banks.
These "demand-supply imbalances" can give rise to non--negligible effects on asset prices and run the risk of compromising our ability to extract the appropriate policy-relevant information from asset price movements.

EXTRACTING INFORMATION FROM ASSET PRICES IN AN UNCERTAIN WORLD

Contents

1. Introduction
2. Forward-looking asset prices and macroeconomic forecasting
3. Extracting information from asset prices
4. Refining our understanding of the role of risk
5. Understanding asset price movements when the textbook model fails
6. Conclusions

References

Extracting information from asset prices in an uncertain world

1. INTRODUCTION

This paper is concerned with how policymakers respond to the presence of uncertainty. Typically, questions addressing this problem are concerned with how policy should optimally be set in the presence of uncertainty about model parameters or about key aspects of the behaviour of economic agents. This paper is concerned with a different aspect to that question, namely how economic agents themselves respond to the presence of uncertainty. This matters because the policymaking problem is usually examined in the context of an economic model where uncertainty is either ignored or subsumed in the constants of the model. This may not lead to serious policy errors if the volatility of the shocks hitting the economy is not changing or if agents' attitude to that uncertainty is constant. But if uncertainty is time-varying then our ability to interpret agents' behaviour and the shocks hitting the economy will potentially be compromised if we ignore that feature of the economic environment.

Asset prices represent a rich and timely source of information about current and future economic fundamentals. They are held because they yield benefits in the future so asset prices change as the markets reassess those benefits. Many of the factors determining asset prices relate to macroeconomic fundamentals, for example prospective real growth or inflation. But some have to do with the investors themselves, for example their attitude to risk or the rules which investing institutions have to follow. And they also provide important information about market participants' perceptions of expected volatility in the environment and their attitudes towards it. This is revealed in observed and expected risk premia, the excess return that investors earn on risky assets. And option prices also allow us to extract information about the expected distribution of asset returns.

The use of asset prices as a source of timely and important information for monetary policymaking is well recognised (see for example Bernanke, 2004, Clews, 2002). This paper examines how information is extracted from financial markets and used to provide policy advice to the Monetary Policy Committee of the Bank of England. It sets out to explain the range of financial market data that are examined to do this, in some cases involving the transformation of raw market data into a form that can be interpreted more easily.

But our ability to "reverse engineer" asset prices to reveal information about the underlying shocks depends on our having a coherent theoretical framework within which to interpret the asset price configuration we observe. So this paper will also describe the modelling toolkit that has been developed to allow us to convert high-frequency readings on financial market prices into policy-relevant guidance on what asset prices are telling us about economic prospects.

These models have importantly allowed us to decompose asset price news into their fundamental determinants; so we have been able to use prices from nominal and index linked bonds to decompose nominal bond yields into real rates and inflation expectations; we can use the Dividend Discount model in conjunction with surveys of corporate earnings expectations to account for equity price movements in terms of changes in interest rates, changes in earnings and the risk premium; we can use the uncovered interest parity condition to check the consistency of observed exchange rate movements with relative movements in interest rates. And we can extract information from option prices to derive market participants' risk-neutral probability distributions of expected future interest rates, equity prices or exchange rates.

One characteristic of these "first generation" models is that in all cases, if the risk premium plays a role at all, it appears as a residual category which acts to reconcile asset price movements with observed fundamentals. So superficially, the "risk premium" is used as a measure of the extent to which the models fail.

But this treatment of the effects of uncertainty on economic behaviour is potentially very misleading. Economic theory predicts that uncertainty will affect consumption and investment decisions, and hence the level of demand. Finance theory tells us that uncertainty will be reflected in asset prices, via risk premia. So we need to understand how risk premia are related to the structure of the macroeconomy.

Indeed, as Cochrane (2006) emphasises,

'The centrepieces of dynamic macroeconomics are the equation of savings to investment, the equation of marginal rates of substitution to marginal rates of transformation, and the allocation of consumption and investment across time and states of nature. Asset markets are the mechanism that does all this equating.'

Asset markets therefore facilitate the flow of funds from savers to investors, and asset returns move to clear these markets. Different types of assets, such as bonds, shares and durable commodities, typically offer very different returns. These returns should, in theory, tell us something about agents' attitudes towards uncertainty

Because of the importance of an understanding of risk premia, recent research in the Bank and elsewhere has attempted to examine the role of risk premia more carefully. Two mutually reinforcing strategies have been taken to do this.

First, a number of approaches have been taken which have attempted to derive empirical estimates of risk premia. In context of "reverse engineering" asset prices, this is a fundamental enterprise since it allows us to identify the genuine expectations (of interest rates, exchange rates etc) of market participants. Without doing this, any attempt to infer underlying movements in fundamentals will be fraught with difficulty.

The second strand of research, equally important, has been to use modern DSGE theory to help us understand clarify how uncertainty impinges on economic behaviour. Due to the technically demanding nature of this problem, this research is in its infancy in the Bank and elsewhere. But it potentially yields us useful insights into how we expect asset prices and risk premia to change in response to changes in the inherent uncertainty of the economic environment or in response in changes in attitudes towards risk.

Of course, even a more carefully articulated approach to understanding risk premia is unlikely to enable us to account for every movement in asset prices in terms of movements in underlying fundamentals. Indeed, it would appear that departures from the textbook paradigm are commonplace in terms of market anecdote relating to, for example, the effect on bond yields of asset-liability management by pension funds or the build-up of dollar reserves by Asian central banks. These "demand-supply imbalances" can give rise to non--negligible effects on asset prices and run the risk of compromising our ability to extract the appropriate policy-relevant information from asset price movements. As yet, we have made little progress in modelling these effects and how they affect our interpretation of asset price movement. But this represents a major challenge to our work programme going forward.

2. FORWARD-LOOKING ASSET PRICES AND MACROECONOMIC FORECASTING

To a macroeconomic forecaster, asset prices represent a rich source of information since. Because asset prices result from the actions of forward-looking investors in well-functioning financial markets, they potentially contain information about market participants' expectations of a range of variables of macroeconomic importance, not least future interest rates.

But asset prices are not the only macroeconomic variables that are determined in a forward-looking manner. Given a range of real and nominal frictions present in the real economy, other economic agents such as firms and households, are also forward-looking and base their decisions today on their expectations of economic conditions in the future. So in principle, the state of the economy today (ie the level of consumption, investment) will be conditioned on a forward-looking view of economic conditions.

In order to try to capture the forward-looking nature of the economy and its central importance in transmission mechanism of monetary policy, the Bank of England's forecast is constructed within the framework of a forward-looking general equilibrium model which explicitly recognises the importance of rational optimising forward-looking behaviour (see Harrison et al., 2005).

2.1 Consistency of conditioning asset price assumptions

Significantly, this interdependence of the information contained in asset prices and the current state of the economy has important implications for macroeconomic forecasting. It means that forecasts ideally need to be constrained such that forward-looking variables are consistent with the future forecast paths of exogenous variables and policy instruments. And since the starting point for the economy is partly determined by the expectations of agents in the economy, it makes sense to ensure that any conditioning assumptions for exogenous variables (such as oil prices), asset prices (such as exchange rates) or policy instruments (such as interest rates) are consistent with the expectations already priced into the market, and hence into the behaviour of agents in the economy.

Partly for this reason, the Bank of England Inflation Report forecast is based on conditioning assumptions which are, as far as possible, consistent with market expectations for oil prices, exchange rates, and most importantly for policy interest rates.

This need to ensure consistency between market expectations and conditioning assumptions has an important practical implication for forecasters who may wish, for good reasons, to base their forecasts on alternative conditioning assumptions which they feel are more plausible or desirable. Any forecast which is conditioned on a different interest rate path than the one embodied in market expectations needs to model explicitly when and how agents respond to the “surprise” that interest rates are following a different path.

Two alternative conditioning assumptions are often used for interest rates:

- Constant interest rate path. This was the assumption used as the conditioning path for the Bank’s main *Inflation Report* forecast up until May 2004[1]. Of course, whichever conditioning path is made, the underlying economic judgements about the behaviour of the economy made within the forecast are no different; the two forecasts represent different angles on the same view, as the Bank’s Governor Mervyn King has explained it. But it can become presentationally awkward to condition on constant interest rates when economic theory would predict reversion to a neutral rate, and when market participants are expecting this;

- A very similar problem emerges for those who advocate that central banks should condition their forecast projections on what they think ought to happen to interest rates (see for example Svensson, 2003). Such an approach is already taken for example by the New Zealand and Norwegian Central Banks (see Norges Bank,

[1] A forecast conditioned on constant rates is still also presented in the Bank’s *Inflation Report*.

2005). It is not the purpose of this paper to argue the pros and cons of that approach, but rather to note that if a forecast is to be conducted (and published) on that basis, it will be necessary for the forecasters to model explicitly how economic agents change their expectations of future policy from what they had originally thought to what policymakers would prefer. And of course, the problem is even more complicated than that because the choice of optimal policy will itself depend on how quickly market participants alter their expectations.

2.2 What is the appropriate market-implied path for interest rates and other asset prices?

Even if a macroeconomic forecast is conditioned on the path for interest rates currently expected by market participants, that still begs the question as to how such expectations should be extracted.

So far the implicit assumption has been made that this is a straightforward task. And indeed, if the pure expectations hypothesis holds, it is. But in practice, the implied rate expectations to be derived from yield curves are likely to be distorted by risk premia and potentially other factors specific to the bond markets concerned. By focussing on bonds issued by governments, the problem of default risk premia is removed, at least for advanced economies. But term premia still remain; this is the premium which bond holders must pay (or receive) for the uncertainty that pertains over the life of the bond due to changes in capital value due to changing interest rates, or changing inflation.

Currently, the conditioning path for interest rates used in the Bank of England forecast does not make an adjustment for term premia over the three year forecast horizon (and beyond 3 years, the interest rate is anyway assumed to revert to a model-based policy rule rather than follow the path implied by market-based forward rates)[2].

But this is an issue which we are working on. Section 3 of this paper will describe how our research agenda at the Bank of England is partly involved with identifying

[2] Some other central banks make simple constant adjustments to account for risk premia, but to our knowledge, none incorporate an explicit time-varying adjustment in their forecast procedures.

and understanding such term premia. This allows us to derive a path for true interest rate expectations held by market participants. And by attempting to estimate risk premia for other asset markets, such as foreign exchange and equity markets, it is similarly possible to forecast asset price trajectories which are consistent with the paths expected in the market.

This does leave one unresolved question, however. If we are able to estimate the interest rate risk premium, does this mean we should discard the risk premium as an annoying distortion? So looking at the stylised paths in chart 1, should we condition the forecast on the green line (as we do now) which includes risk premia, or the pink line, which strips them out. In the context of a modelling framework which effectively treats risk as a nuisance constant, the path given by pure expectations seems the obvious choice. But it is important to remember that risk premia are a manifestation of economic agents' response to uncertainty. And as such, the presence of changing perceptions of uncertainty should be affecting our view of macroeconomic prospects, so we would be wrong simply to throw away the information in risk premia. Our attempts to articulate these influences more rigorously are described in Section 4.

Chart 1: What is the appropriate conditioning path for interest rates: with or without risk premia?

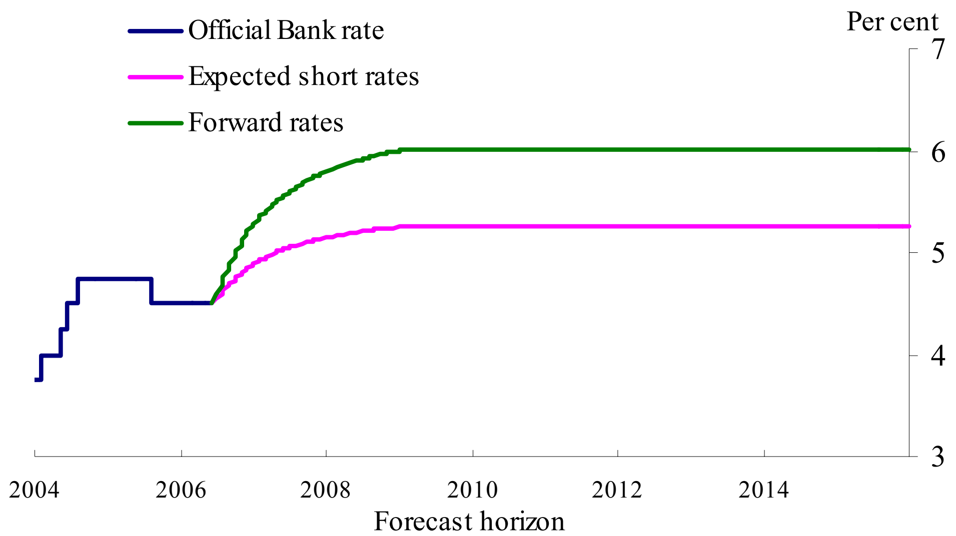

3. EXTRACTING INFORMATION FROM ASSET PRICES

Asset prices can change for a number of different reasons which may have different implications for the inflation outlook.

- Market participants' best assessment of the future may have changed because their view of the economy has changed

- Or because their view of the MPC has changed

- Or market participants' view of the future may have become more or less diffuse (risky)

- Or they may be more or less comfortable about bearing risk (there is evidence that the strength of international influences on domestic financial markets varies over time, being especially high at times of financial market stress[3].

The Monetary Policy Committee may be interested in all of these elements, and may wish to take account of them in different ways. But unscrambling the different elements from observed asset prices is rarely easy.

As a consequence, it is important that asset price information is assessed in the context of a well-defined theoretical framework. This section of the paper describes the range of asset price information that is presented to the MPC and explains some of the simple (and sometimes not-so-simple) tools that have been used to interpret asset price movements.

For each asset class presented (interest rates, equities, exchange rates and derivatives), the basic tools are first presented and explained. In most cases, the presence of risk premia in the basic models is noted as a complication, or indeed as a measure of our inability to explain asset price movements, rather than as a reflection of investors potentially changing attitudes to risk and uncertainty. So for each asset class, the paper will then explain recent research at the Bank of England which aims to quantify risk premia more accurately. Our attempts to understand and model the theoretical determinants of risk premia will be left to section 4.

[3] For the case of bond markets, see Clare, A and Lekkos, I (2000).

3.1 Interest rates

The near-term market expectations of interest rate changes provide the MPC with important information about the extent to which planned policy changes are anticipated. This matters because the effects of unanticipated policy changes have different economic effects to anticipated moves.

Short-term interest rate expectations

For short-term interest rate expectations over the year ahead, overnight interest rate swaps (in the UK, euro area and US) can be used to derive near-term expectations of policy rate changes at particular meetings (see chart 2a). And for expectations slightly further out, changes in interest rate futures markets can provide useful information, for example, on how particular items of economic news have impinged on rate expectations; chart 2b illustrates[4].

Chart 2a: SONIA step chart for UK near-term interest rate expectations

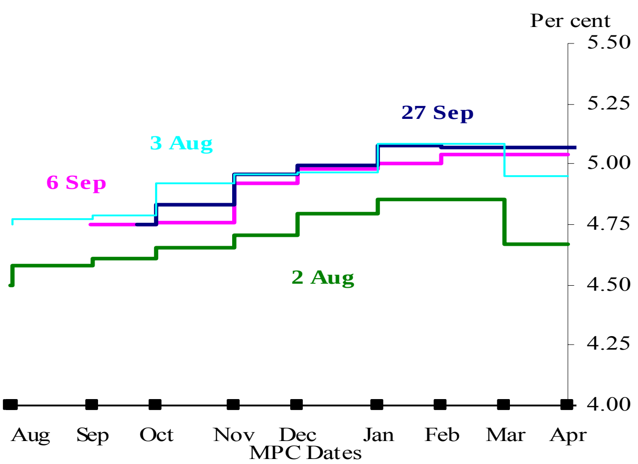

Chart 2b: Using international interest rate futures to assess the effect of news

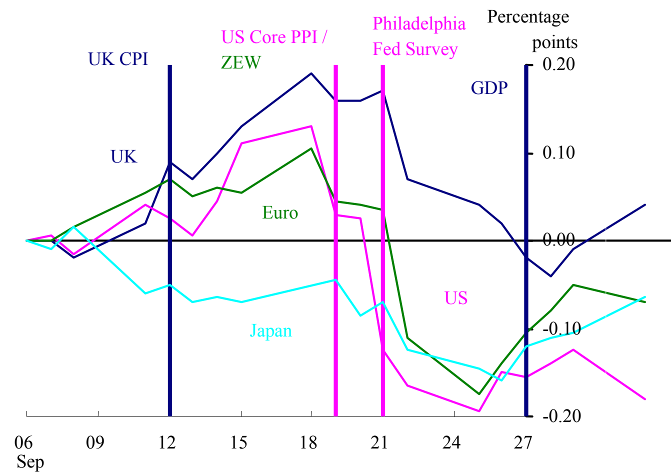

For interest rate expectations at longer maturities, it is necessary to extract market expectations from bond prices. These are potentially a rich source of information about market participants' expectations of future monetary policy and the shocks that are expected to impinge on the economy. In the UK context, the MPC sets a short-term nominal interest rate – the Bank Rate - which directly influences other short-

[4] Since interest rate futures relate to interbank interest rates, these will be subject to a credit risk and so require a credit risk adjustment before any interpretation is put on the implied levels in terms of expected future policy rates.

term rates in sterling markets. But expectations about rates to be set by the MPC in the future also have an immediate impact on yields on long-term bonds held over a number of periods[5]. For example, if short-term interest rates are expected to rise in the future, then long-term interest rates will reflect that today since investors always have the option of instead rolling over a series of short term contracts. So by comparing the return on long-bonds of different maturities, we can calculate what the series of short-term interest rates would need to be to earn the same return as the long-term interest rate. This series of implied short rates are known as forward rates (and later on we will discuss the extent to which they correspond to genuine expected future policy rates).

But before it's possible to extract policy-relevant information from bond yields in this way, it is important to recognise that a considerable amount of effort needs to be put in to transform the raw bond price data into the yield curves that we typically find useful. So yield curves need to be estimated; they can't just be read from FT or Bloomberg!

Prices of nominal bonds provide information about redemption yields of bonds of different maturities. In the case of pure bonds, where the borrower is required to repay the principal amount borrowed at the redemption date, then it is straightforward to derive implied short-run interest rate expectations by comparing the yields on bonds of adjacent maturities (so for example the expected one-year[6] interest rate five years ahead could be derived by comparing the yield on a four and five year bond). But conventional bonds often mature at relatively infrequent intervals and have periodic coupon payment, so before a set of comparable spot and forward yields can be derived from bond prices, it is necessary to estimate continuous zero coupon curves. A range of methods are available to do this. The methodology used at the Bank of England is based on a VRP (variable roughness penalty) method rather than so-called Svensson curves; for further details, see Anderson and Sleath (1999). Chart 3 illustrates how the eventual estimated spot curve is related to the original observed yields.

---

[5] The n-period long rate, or spot interest rate for *n* years, refers to the interest rate applicable today ('spot') on an *n* year nominal loan. It is the average rate at which an individual nominal cash flow on some future date is discounted to determine its present value.

Chart 3: Moving from observed yields to an estimated spot yield curve

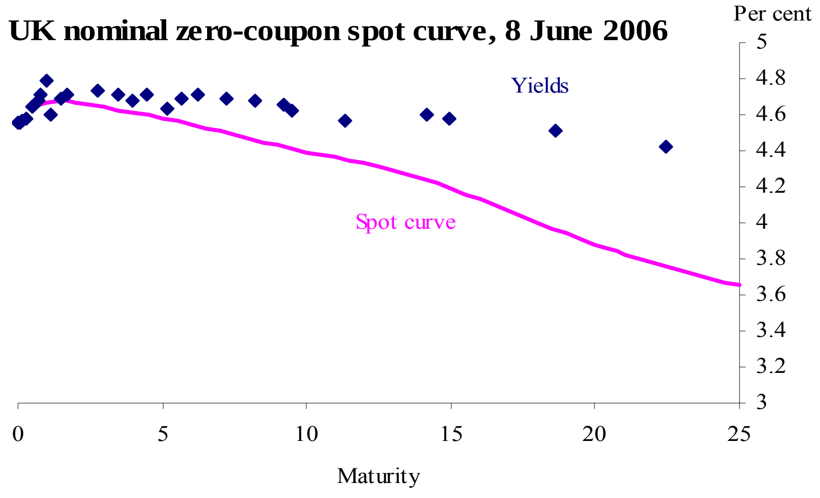

Yield curves are estimated at the Bank of England for a range of countries on a daily basis. The main curves used in the briefing of the MPC are those for the UK, US, euro area and occasionally Japan.

Nominal yield curves

Two kinds of nominal yield curve are estimated for the UK. One set is based on yields on UK government bonds and on yields in the general collateral repo market; this is known as the government curve. The other set is based on sterling interbank rates (LIBOR) and on instruments related to LIBOR (short sterling futures contracts, forward rate agreements and LIBOR-related interest rate swaps); this is known as the Bank liability curve (BLC)[7]. Both these yield curves are published daily on the Bank's external website. Daily calculations are also made of the nominal yield curves for the US and euro area.

Chart 4a below shows a typical yield curve comparison presented to the MPC.

---

[6] This is known as the "tenor" of the forward rate. Similarly it is possible to derive forward rates at shorter tenors, where in the limit, we can focus on the instantaneous forward rate.
[7] The way in which the methodology is adapted for the commercial bank liability curves is described in Brooke, Cooper and Scholtes (2000).

Chart 4a: Nominal yield curves for the UK, US and euro area

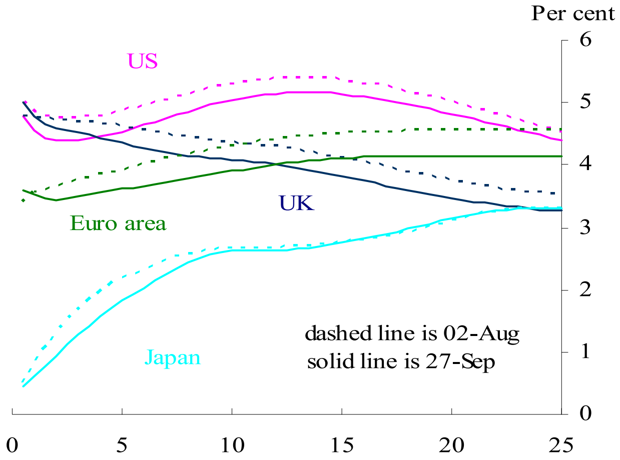

Chart 4b: The UK adjusted BLC curve (used as interest rate conditioning path for forecast)

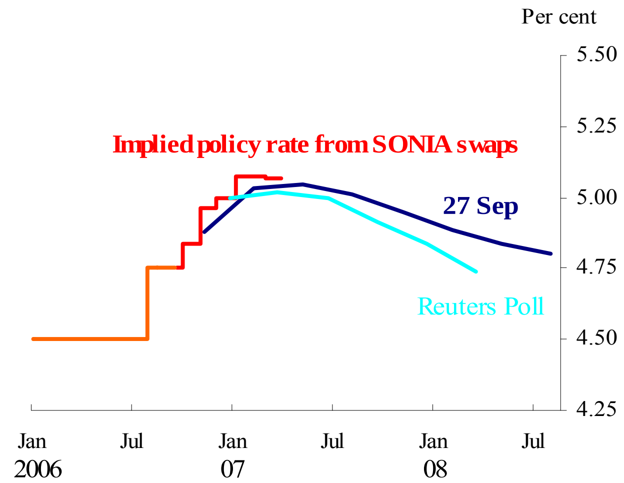

In principle, because government bonds are free of default risk, forward rates from the government curve provide a clean read on expectations of future policy rates at least at short horizons where term premia are not large. In practice, especially in recent years, demand-supply imbalances have implied that the government curve has provided a distorted view of likely short-term interest rate expectations in the UK. As a result, we have tended to rely on the UK BLC for this purpose. Because this curve is based on commercial bank liabilities, the resulting interest rate trajectory will be distorted upwards by an effect from credit risk (associated with possible bank default). An adjustment is made to the BLC curve to remove this credit risk effect (where the credit risk adjustment is calibrated by comparing the BLC curve with the policy rate and the gilt repo rate at the short end of the curve and assuming that the credit risk adjustment is constant for maturities beyond one year). This so-called adjusted BLC curve is the curve used for the first three years of the interest rate market conditioning path in the Bank forecast (see chart 4b above).

Decomposing the nominal curve into real rates and inflation breakevens

The return on a nominal bond can be decomposed into two components: a real rate of return and a compensation for the erosion of purchasing power arising from inflation. For conventional government nominal zero coupon bonds, such as those in the example above, the nominal return is certain (provided it is held to maturity) but the real return is not (because inflation is uncertain). An index-linked zero coupon bond would have its value linked to movements in a suitable price index to prevent inflation eroding its purchasing power (so its 'real value' is protected). For such a zero coupon bond the real return would be certain if the bond were held to maturity. A real debt market provides information on the *ex ante* real interest rates faced by borrowers and lenders who want to avoid the effects of inflation which allows us to calculate real spot and forward rates analogous to the nominal spot and forward rates described above.

We have seen that the index-linked gilt market allows us to obtain real interest rates and the conventional gilt market allows us to obtain nominal interest rates[8]. These nominal rates embody the real interest rate plus a compensation for the erosion of the purchasing power of this investment by inflation. The Bank uses this decomposition (commonly known as the Fisher relationship) and the real and nominal yield curves to calculate the implied inflation rate factored in to nominal interest rates. This is often interpreted as a measure of inflation expectations, although some care is required in doing so. As with nominal and real interest rates, we can think of 'spot' implied inflation rates as the average rate of inflation expected to rule over a given period. Similarly forward implied inflation rates can be interpreted as the rate of inflation expected to rule over a given period which begins at some future date. In the limit, we can calculate instantaneous forward implied inflation rates just as with real and nominal rates.

We can only derive real yield curves (ie spot and forward real yield curves) in those countries where index-linked bonds are issued, as detailed below:

- It has been possible to derive such curves for the UK since 1985.
- The development of the TIPs market in the US mean that we can also produce real yield curves for the US since 2000.
- Some euro area countries do issue index-linked bonds but not enough to derive a full real yield curve across a wide range of maturities. But we now exploit information from inflation derivatives markets which allows us to derive inflation

[8] Strictly speaking, we also need to use the nominal curve to drive real rates from index-linked bonds because we have to adjust for the imperfect indexation due to the indexation lag.

expectations for the euro area (and indeed for the UK and US across maturities). For more detail, see Hurd and Relleen (2006).

Charts 5a and b shows the resulting international comparisons of UK, US and euro area real interest rates and inflation breakevens from these markets, as typically presented to the MPC.

Chart 5a: Real interest rates for the UK, US and euro area[9]

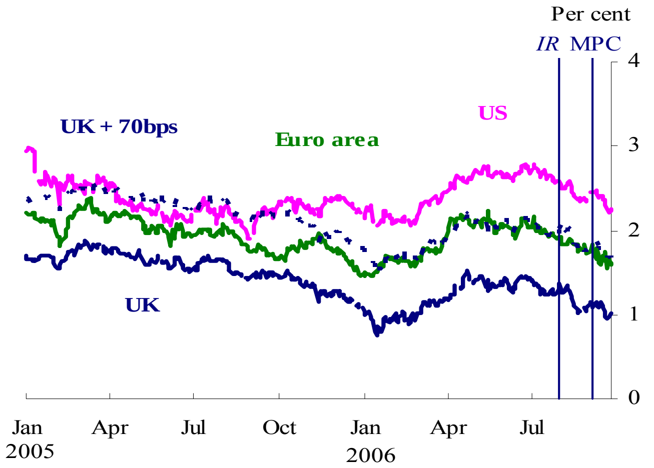

Chart 5b: Inflation breakevens for the UK, US and euro area

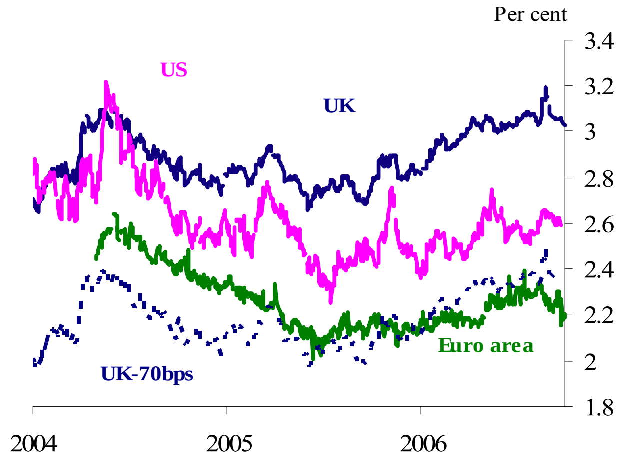

Interpreting yield curve movements

While our ability to decompose yield curves into real rates and inflation expectations is vital for understanding why yield curves are moving, this is only a first step. If the expectations hypothesis held, then we might expect the derived real rates and inflation breakevens to be giving us a clean read on expected future real rates and expected future inflation. But in the presence of risk and uncertainty, life may not be so simple. If investors are averse to risk, as empirical evidence suggests they are, then observed real yields will actually comprise forward real yields plus a term premium. And similarly, inflation breakevens will include an inflation risk premium.

Early approaches to modelling the yield curve tended to adopt a finance-based approach which explained the yield curve in terms of itself, a long rate and a spread

[9] UK real rates are adjusted upwards to account for the fact that UK index-linked bonds are indexed to RPI inflation rather than a CPI-based measure. In the UK, the steady-state wedge between RPI and CPI inflation is estimated to be around 70 basis points.

factor (for an example of earlier Bank work in this area, see Steeley, 1997). But more recent research has attempted to combine the insights of this traditional 'finance' modelling approaches with models that also relate yield curve movements to macro-factors. Our current and proposed work is trying to draw on the insights of this literature.

First, we have completed a project which estimates an arbitrage-free affine macro-factor model of the UK nominal yield curve (see Lildholdt *et al.*, 2006). This has yielded a number of insights into the main macro-factors driving changes in the UK nominal yield curve, in particular finding that the inflation target plays a key role in explaining the evolution of long forward rates. The model also provides us with a way of decomposing the yield curve into interest rate expectations, risk premia and convexity effects. Charts 6(a) and (b) illustrate the key results.

Chart 6a: Factor evolutions	**Chart 6b: Forward term premia**
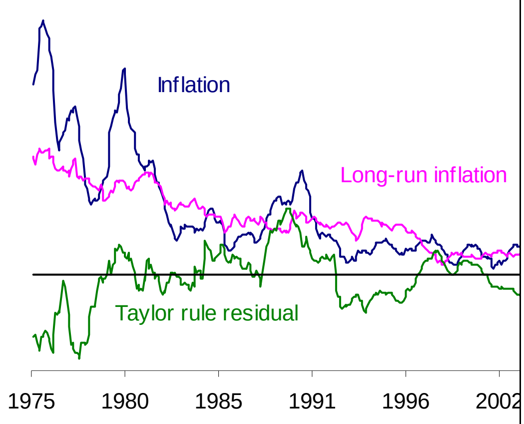	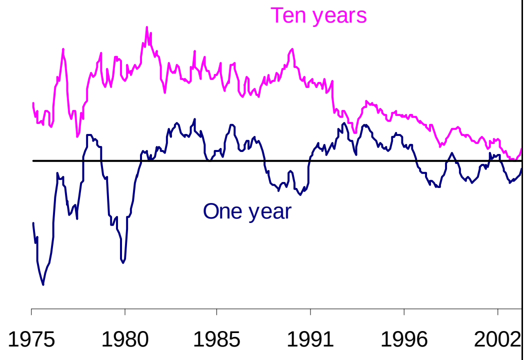

Second, we have estimated separate affine latent factor models of the nominal and real term structures. These 'finance' models have a more flexible structure than either macro-factor models or GE models, making them more robust to model misspecification. The assumptions in the model are that there are no arbitrage opportunities and that yields are driven by a small number of latent factors. Given these assumptions, the models provide a theoretically consistent way of decomposing forward interest into expectations of future short rates and forward term premia.

Our results using this approach on the real term structure have helped inform our analysis of the so-called bond market "conundrum" of low long real rate. Some recent widely publicised yield curve modelling work by the Fed (see Kim and Wright, 2005) has attempted to decompose yield curve movements into those caused by expected real policy rates and those caused by risk premia. Their conclusions, which have been cited in speeches by Fed officials, are that expected real policy rates changed little over the 1990s and that the fall in real rates can be attributed to lower risk premia. Chart 7a captures their main result where over the 90s, the 10 year forward rate (the blue line), breaks down into the underlying policy rate (pink) ,which is largely unchanged, and the risk premium (the green line) which has fallen.

In parallel, we've conducted similar work for the UK, using information from our index linked bond market to model real yields directly, rather than backing out implied real yields using a model of inflation, as the Fed researchers are obliged to do (given the absence of a long enough back run of TIPS data).

Chart 7a: Decomposing the US 10-year real forward rate

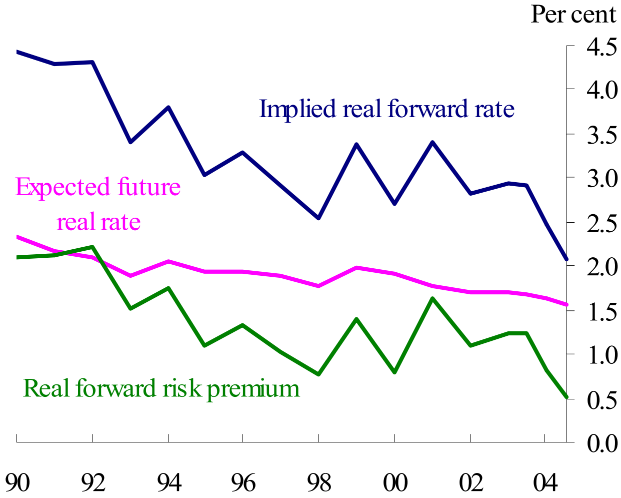

Chart 7b: Decomposing the UK 10-year real forward rate

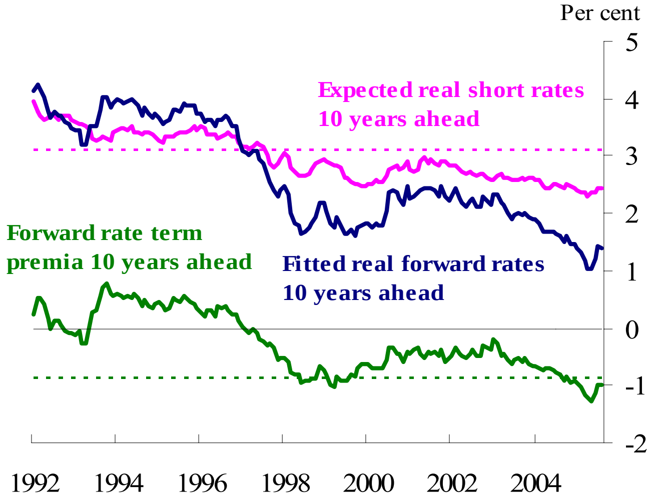

Chart 7b shows our corresponding results for the UK. The results are similar to those found for the US, namely that most of the fall in forward rates is associated with real "risk premia". The model results show expected future policy rates remain close the model's implied long-run equilibrium, providing little reason to expect this component of the real rate to change much. But given that real term premia are currently negative, and about 100 basis points below the level the model expects them to reach in the long-run, there is a clear upside risk to long real rates going forward.

Of course, one has to careful in interpreting these results too literally. These models are reduced-form and do not explain the fundamental causes for the fall in real risk premia. They predict that underlying real rates and term premia will always revert to a constant value over the very long run. And there is no explicit account for market segmentation effects. But even on that basis, we still find these results interesting.

Following on from this work, we are now attempting to model the real and nominal term structures jointly using a similar framework, but allowing for both unobservable and observable factors. This joint model should enable us to model the inflation risk premium and to generate market-derived expectations of future real rates and inflation rates. The ultimate aim of this work is to be able to decompose nominal forward rates into expected real rates and inflation, real and inflation risk premia and real and inflation convexity effects.

3.2 Equity prices

A range of equity market indices are presented to the MPC, typically covering the FTSE All Share for the UK, S & P 500 for the US and Euro Stoxx for the euro area; see chart 8.

Chart 8: International equity price movements

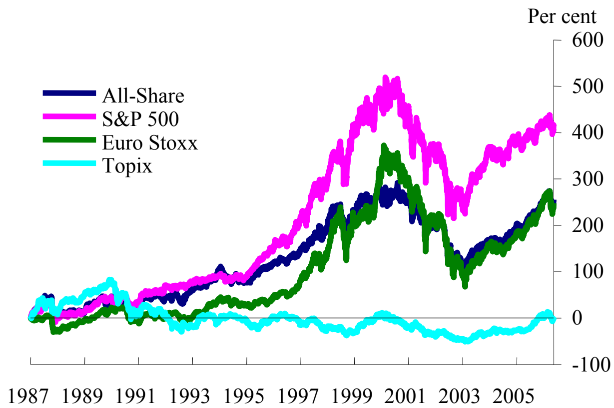

Interpreting movements in the valuation of equities is notoriously difficult. Movements in share prices depend on prospects for the relevant companies. Investors and market analysts scrutinise the outlook for individual companies so that share prices are in principle summarising a mass of detailed information that could well be of interest to the MPC – not only on the economy as a whole but also on particular sectors. For example, movements in equity prices are likely to reflect changing perceptions of future productivity growth. For markets, productivity growth matters because it affects future corporate earnings and dividends. To policy makers it matters because it helps determine the rate at which the economy can grow without putting upward pressure on inflation. But equity values depend not only on future streams of earnings or dividends but also on the rate at which they are discounted, including any allowance for the risk involved. If equities come to be seen as less risky that would raise share prices, even if there was no change in the view taken of future earnings or dividends (see Panigirtzoglou and Scammell, 2002).

In practice, care is needed in weighing the significance of UK asset price movements for domestic prospects. Many large UK companies have interests overseas, and movements in UK equity prices may reflect prospects for overseas subsidiaries as well as domestic activity. Moreover there is evidence that the strength of international influences on domestic financial markets varies over time, being especially high at times of financial market stress (see for example, Clare and Lekkos, 2000).

The Dividend discount model

The dividend discount model (DDM) is the workhorse of much asset pricing work especially in the context of the extraction of information from equity prices. By articulating equity price movements in a well-defined theoretical framework, it facilitates a systematic decomposition of movements in equity prices into their fundamental macroeconomic drivers.

In fact, the dividend discount model (DDM) can be applied to any asset. It states that the price of an asset (P_t) should equal the discounted expected sum of any future transfers (D_t) to the asset holder:

$$P_t = \mathrm{E}_t \sum_{i=1}^{\infty} \frac{D_{t+i}}{(1+R_{t+i})} \tag{1}$$

where the discount factor R_{t+i} reflects the risk free rate (r^f) and a risk premium (ρ) which ensures that agents are appropriately compensated for the uncertainty surrounding the flow of future cashflows in relation to their income.

To make this operational, we need to be able to monitor available information about

- expected future dividends; we can do this in a number of ways. In the simple 1-stage DDM (also known as the Gordon growth model) we can make a simple assumption about the long run growth of dividends. This can either be determined by assuming that dividends grow in line with money GDP, or more sophisticatedly by calculating the expected steady-state growth rate of dividends (which will depend on the retention ratio and the return on any reinvested earnings[10]). In the three-stage DDM, additional information is used about the cyclical behaviour of corporate earnings before they reach their steady-state growth path (as calculated above). This is derived from surveys of institutional investors (for more details, see Panigirtzoglou and Scammell, 2002).

- expected future real interest rates; these are derived from the forward yield curve (in practice out to the ten-year horizon).

Chart 9a: DDM decomposition of the UK FTSE All share index (since August Inflation Report and September pre-MPC)

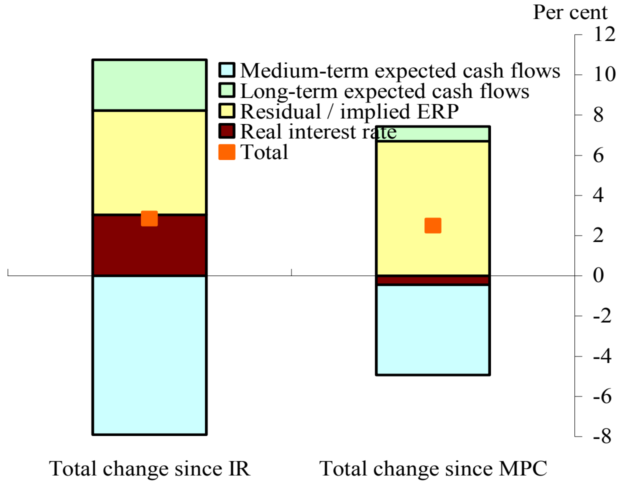

Chart 9b:

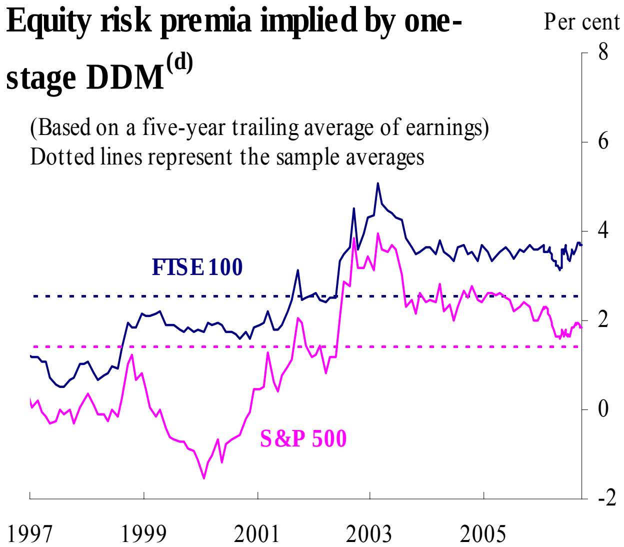

Given these assumptions, the equity risk premium is the remaining unknown to account for observed equity price movements. Chart 9a shows an example of how FTSE all share equity price movements can be decomposed using this model.

[10] This will be given by the return on equity multiplied by the retention ratio.

This model is also informative for gauging the valuation of equity markets relative to some long-run average. Although the risk premium can fluctuate in the short-run, perhaps due to cyclical risk aversion, we believe it will revert to some long-run 'sustainable' level which we proxy with the sample average. So the current position of the risk premium relative to this long-run average is an indication of whether the asset is 'fairly' valued.. For example, chart 9b shows how the UK and US equity risk premium[11] implied by the 1-stage DDM has moved relative to its historical average.

The role of risk premia in equity valuations

In practice, however, backing out the risk premium in this fairly naïve way is problematic because what we obtain is a residual. Our measure of the risk premium thus conflates the true risk premium with the effects of any problems with the model or the data. So in recent work, we have attempted to think more carefully about the role of risk premia, in particular in the context of the role of low interest rates in driving equity valuations.

One of the risks identified in the November 2005 Inflation Report was that a rise in long-term real interest rates from their historically low levels could trigger a fall in asset prices. To examine that risk more clearly, it is important to understand why real interest rates might be so low and under which circumstances a sudden increase would lead to a substantial fall in other asset valuations.

To illustrate this question, Chart 10 shows the relative valuation of the FTSE All-Share index since 1982. Focussing on the end-2005 readings, it shows that based on current real long-term interest rates and an implied equity risk premium equal to its long run average since 1982, UK equities are about 14% undervalued.[12] But real long-term rates are unusually low at the moment, around 1% for the 10-year real spot rate. If long-term real interest rates were to return to more "normal" levels, say about 3%

[11] This is calculated on the assumption that the equity risk premium is expected to remain constant over the future life of the asset.
[12] Given the margins of error with these calculations, it is important not to get too concerned about the precise numbers associated with a particular valuation measure.

(the average of the 10-year real rate since 1982), then UK equities would be almost 25% overvalued.[13]

Chart 10: One-stage DDM valuation: FTSE All-share

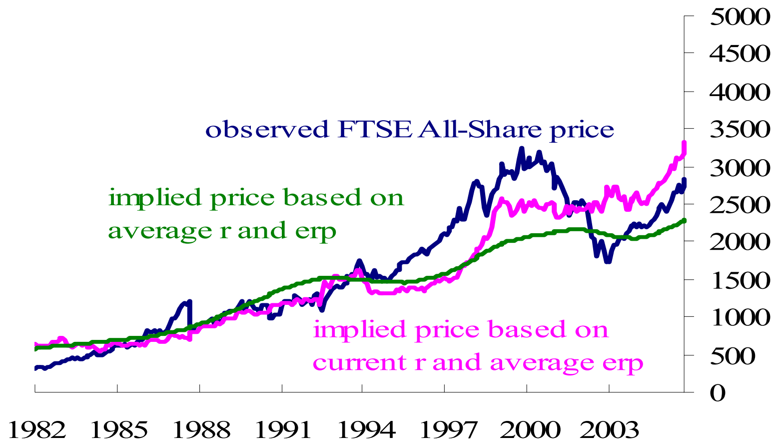

In fact, this isn't just an UK equity markets phenomenon. On the same basis US equities would start off roughly fairly valued, but at average rates would be a massive 48% overvalued (again based on end-2005 calculations). Moreover, this type of argument could be applied to a wide range of domestic and global assets such as housing, fine art and gold, for example.

This simple illustration clearly illustrates that it is important to think carefully about deriving measures of equity market valuation on current market interest rates, since current long rates used in the DDM model will themselves embody a term premium which may itself be unsustainably low.

Ideally, to address this problem more formally, we would like to be able to estimate joint models of the equity risk premium and the term premium embodied in the yield curve. Indeed, theoretically, these two concepts are closely related; the equity risk premium comprises the sum of the term premium and an additional premium for payout uncertainty associated with the uncertain of the dividend stream (see for example Jermann, 1998, and Beckaert et al., 2005).

[13] Of course, this thought experiment importantly requires holding all other factors equal such as the equity risk premium and future corporate earnings. As this note goes on to explain, different

3.3 Exchange rates

Exchange rates are the price of one currency in terms of another. So there is a wide range of factors in this country and abroad that in principle bear on sterling's exchange rates against other currencies. But quantifying the effect of each has often proved very hard. For example, the largest movement in sterling in the last decade occurred in the second half of 1996 (ie just before the establishment of the MPC) and was analysed by Bank staff in the *Inflation Report* of February 1997. The factors considered were monetary and fiscal policy at home and abroad, portfolio shifts associated with the prospect of EMU, movements in the oil price, possible shifts in the demand for UK goods in world markets, and possible improvements in productivity in UK industries producing internationally tradable goods and services. The implications for inflation of each of these factors could be quite different, but it was hard to know how important each of them was. Exchange rate movements since the establishment of the MPC have been less extreme.

The UIP decomposition

As with equity prices, it is important to be able to analyse exchange rate movements in the context of a well-understood theoretical framework. The workhorse model of exchange rate movements is the UIP decomposition (for a comprehensive decription of this approach, see Brigden et al., (1997). This attempts to account for exchange rate movements conditional on what has happened to relative interest rate movements according to the uncovered interest parity condition.

i.e. $$s_t = E_t s_{t+1} + \left(i - i^*\right)_t + \phi_t$$

where s_t is the nominal exchange rate, i_t the domestic interest rate, i_t^* the domestic interest rate and ϕ_t the foreign exchange risk premium.

Or ignoring the risk premium term (for now),

i.e. $$s_t = E_t s_{t+1} + \left(i - i^*\right)_t$$

---

explanations for the fall in real rates may well be associated with changes in these other factors

If we re-write the UIP expression in real terms[14] to leave us with a UIP condition in real terms;

i.e.

$$s_t + p_t - p_t^* = E_t(s_{t+1} + p_{t+1} - p_{t+1}^*) + \left((i_t - E[p_{t+1}] + p_t) - (i_t^* - E[p_{t+1}^*] + p_t^*)\right)$$

which can be re-written as

$$er_t = E_t er_{t+1} + \left(r - r^*\right)_t$$

where er_t and r_t are real exchange rates and interest rates respectively.

Now we can integrate the expression forwards;

i.e. $$er_t = E_t er_{t+n} + \sum_{k=1}^{n} E_t\left(r_{t+k-1} - r^*{}_{t+k-1}\right)$$

This gives an expression which relates the current real exchange rate to the real exchange rate at some arbitrarily distant point in the future, plus a term in cumulated expected real interest rate differentials.

Of course, when we are considering moves in the exchange rate between adjacent dates, we need to consider this expression in terms of changes in the right-hand-side variables between those dates. And we can use this expression to derive the result that the jump in the nominal exchange rate will equal the change in the cumulated real interest rate differential. This requires the following assumptions.

- the equilibrium exchange rate hasn't moved between the dates under consideration;

- the cumulated real differential is bounded, ie changes in real rate differentials converge to zero;

- the price level at time zero doesn't move; this allows us to assume that the jump in the nominal and real exchange rate is identical.

[14] This is done by adding and subtracting domestic and foreign inflation from both sides of the equation.

In the absence of information about real interest rate differentials, we need to make an assumption about how the observed movements in nominal interest rates is accounted for by movements in real rates. Typically, in presenting information to the MPC, we make two assumptions, both of which assume that the change in real yield differentials converge to zero:

- *monetary news*; At the very short horizon, relative interest rate movements is assumed to be entirely real, this declines linearly until 6 years after which it is entirely nominal (see chart 11a for a stylised representation).
- *nominal interest rate news*; All relative interest rate movements out to the ten year horizon reflect real interest rates; beyond that real interest rate differentials are zero.

Charts 11 gives an example of an actual change in relative interest rates, and shows the assumptions made in the monetary news calculation; the interest news calculation would assume that all movements out to ten years were driven by real rates.

Chart 11a: Stylised chart of assumptions made in monetary news calculation

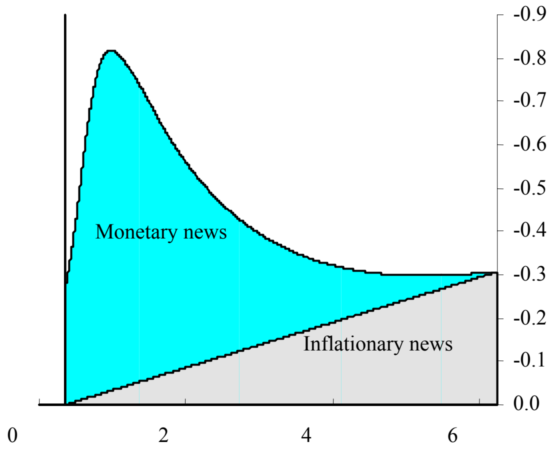

Chart 11b:UIP decomposition of changes in £ERI since August 2006 Inflation Report

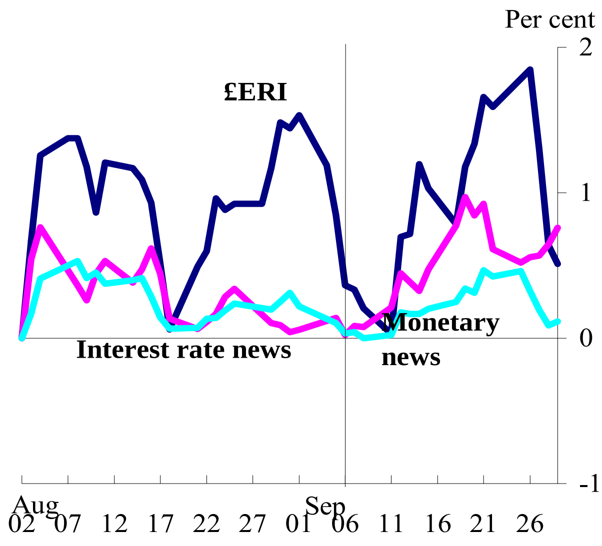

Of course, in practice, the resulting UIP decomposition is usually unable to account for the observed exchange rate moves. According to this framework, the discrepancy is either explained

- by moves in the equilibrium exchange rate; or

- by the fact that the riskless UIP framework is wrong. So this encompasses (a) the possibility that the UIP expression should be augmented by risk considerations; or (b) the UIP framework is not valid, for example due to segmented markets etc.

So yet again, as with equities and interest rates, the workhorse tool for analysing exchange rates ends up treating risk premium considerations as a residual measure.

Incorporating real yields and risk premia in the UIP decomposition

One obvious problem with the way the UIP decomposition is currently implemented is its reliance on an arbitrary assumption about the movement of relative interest rates. In principle, given the availability of market-based expectations of real interest rates, it might seem to be straightforward to make this improvement. In practice, however, there is a complication. The logic of the decomposition is that the forward accumulation of real interest rate differentials needs to be bounded (ie real interest rates are assumed to converge together, consistent with the notion that the real exchange rate should be constant in the long run). In practice, observed real interest rate differentials sometimes move in parallel. Theoretically, we can rationalise this in terms of movements in relative risk premia. Moreover, it is possible to show theoretically that movements in FX risk premia are closely related to movements in relative term premia (see Backus et al., 2001) So as with equity risk premia, our research goal is to derive empirical estimates of the FX risk premium, consistent with our term premium estimates from the yield curve, and use these to modify our UIP decomposition to incorporate risk premia. Preliminary work along these lines includes Backus et al., 1991, and more recently within the Bank, Benati (2006).

3.4 Derivatives

Some financial instruments, options, are particularly related to risk. They are designed to put a value on the risk of future movements in the price of an underlying asset, which is often a financial instrument such as an equity or foreign exchange, but can also be a physical commodity such as oil.(15) From options prices we can derive measures of the expected volatility of certain asset prices. And on further assumptions, we can infer "implied probability distribution functions (pdfs)" of future

(15) Options give the right, but not the obligation, to buy or sell the underlying asset at a point in the future at a price set now (the strike price). An option to buy only has value if there is a chance that the underlying asset will be worth more than the strike price when the option comes to be exercised. So by examining prices of options at different strike prices we can form a view on the probability that the price of the underlying asset will be in different ranges.

asset prices, rather like the fancharts which express the MPC's own uncertainty about future inflation or about GDP growth.[16] So the width of the pdf will reflect uncertainty about future asset prices. And the extent to which the pdf is asymmetric can potentially tell us about market views on the relative risks that future asset prices will be higher or lower, the so-called 'balance of risks'.

Information about the shape of implied pdfs for different asset prices forms part of the information set regularly examined by the Bank in pursuing its two Core Purposes of monetary stability and financial stability. It is primarily of use in helping policymakers to understand market expectations about a range of future asset prices — and, by extension, perhaps the economy. For monetary stability, interest rate probability distributions implied by option prices are one way of assessing market views about risks around the path of expected future interest rates. Such views could reflect market uncertainty about the monetary policy reaction function or about the nature of exogenous risks facing future interest rates and the economy. Turning to financial stability, information from option prices could be useful in monitoring and assessing potential risks to the financial system. For example, concentrations of probability in the tails of the probability distributions for future asset prices may indicate growing perceptions of a risk of unusual movements in asset prices.

Generating option-implied pdfs in the Bank of England

Specifically, in the Bank we derive daily pdfs for equity indices, interest rates, exchange rates and commodities; summary indicators of these pdfs for the FTSE 100 are short sterling interest rates are provided daily on the Bank's public website.
For example, Chart 12a shows the implied volatility of US and UK equity prices derived from equity options; Chart 12b shows a probability distribution of expected interest rates at a 3 month horizon, as routinely showed in the Bank's *Inflation Report.*

---

[16] See **Clews, R, Panigirtzoglou, N and Proudman, J (2000)**, 'Recent developments in extracting information from options markets', *Bank of England Quarterly Bulletin*, February, pages 50-60.

Chart 12a: FTSE 100 3month forward implied volatility

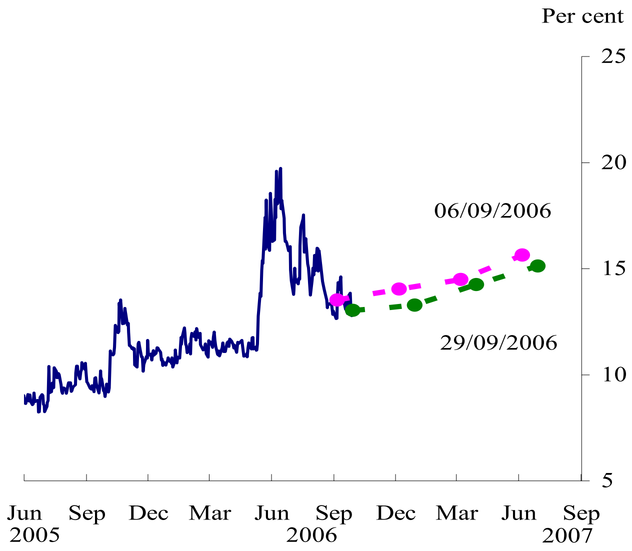

Chart 12b: Interest rate PDF derived from Libor futures

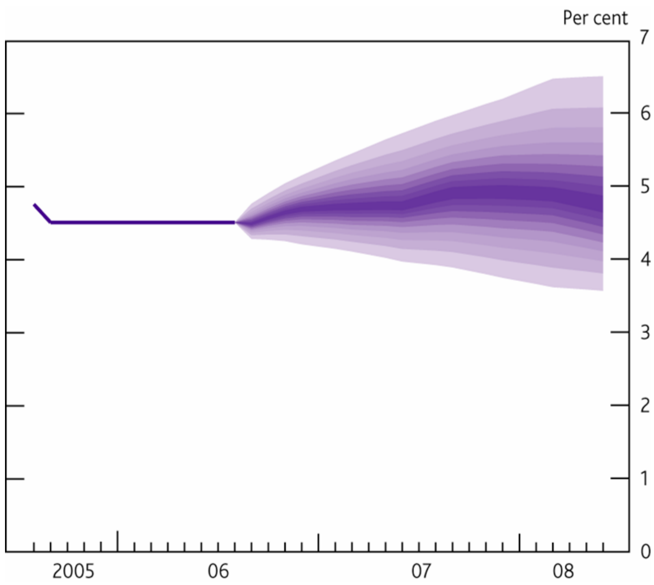

Technical considerations issues with pdfs

Formally, option prices can be used to infer a set of risk-neutral probabilities attached by financial markets to various future asset price levels. In the jargon this is referred to as an option-implied probability density function (pdf) for the price of the underlying asset in the future. Technically, it is constructed using the non-parametric method of Breedon and Litzenberger (1978). A cubic spline is fitted through the volatility smile of daily settlement prices. The spline is then transformed back into a price function, evaluated over a fine grid; taking the second derivative gives the PDF. We also use this spline to estimate the implied volatility of an option that is exactly at-the-money.

Options contracts traded on financial futures exchanges, such as LIFFE, have fixed expiry dates corresponding to the maturity of the underlying futures contracts: March, June, September and December. This feature can make comparing pdfs over time difficult. This is because the degree of uncertainty about the price of the underlying futures contract at the expiry date of the option naturally decreases as the expiry date approaches. So the implied volatilities and variances of the pdfs that we estimate diminish over time, without any real change in the degree of uncertainty about the asset. Normally the implied volatility of each contract drifts downwards through the operation of this 'time-to-maturity' effect. But volatilities can also be shocked by some external event. To discern more clearly such underlying changes we need a method for stripping out the 'time-to-maturity' effect. Our method for doing this is based on—and is consistent with— the non-parametric technique discussed above. There we interpolated across the implied volatilities of options with different deltas but with the same maturity. Here we interpolate across the implied volatilities of contracts with the same delta but with different maturities. In fact, we can construct a surface of implied volatility. The surface is estimated from the implied volatilities from contracts on all available deltas and maturities. The implied volatility smile of a constant-horizon pdf can be thought of as a cross-section of the surface at a particular date.

Option implied joint-PDFs – the sterling ERI

As there is no actively traded market in options on the sterling ERI, the risks to the future value of the index cannot be inferred directly. An indirect approach is to model the probability distribution of a 'simplified' sterling ERI based solely on the prices of options on euro-sterling and dollar-sterling, which are the key bilateral exchange rates in the sterling ERI.

A statistical tool (known as a copula function) can be used to map the euro-sterling and dollar-sterling implied distributions onto a joint distribution (see chart 13a; for more details, see Hurd et al, 2006). From the joint distribution, an estimate of the implied probability distribution for the simplified sterling ERI can be constructed. More specifically, using weights of 0.7 for euro-sterling and 0.3 for dollar-sterling, it is possible to back out an implied probability distribution for the sterling ERI (see chart 13b). Given the probability distribution for the sterling ERI, it is possible to construct synthetic measures of implied volatility and risk-reversal statistics. Another use of the joint distribution is to calculate conditional probabilities. For example, it is possible to construct an implied distribution for the sterling ERI over the next twelve months given an assumed change in the euro-dollar exchange rate. At face value, the distributions indicate that option market participants perceived that even relatively large bilateral movements in the value of the dollar against the euro would tend have a relatively modest impact on the probability distribution of the ('simplified') sterling ERI.

Chart 13a: Bivariate distribution of £-$ v £-€ implied by copula function

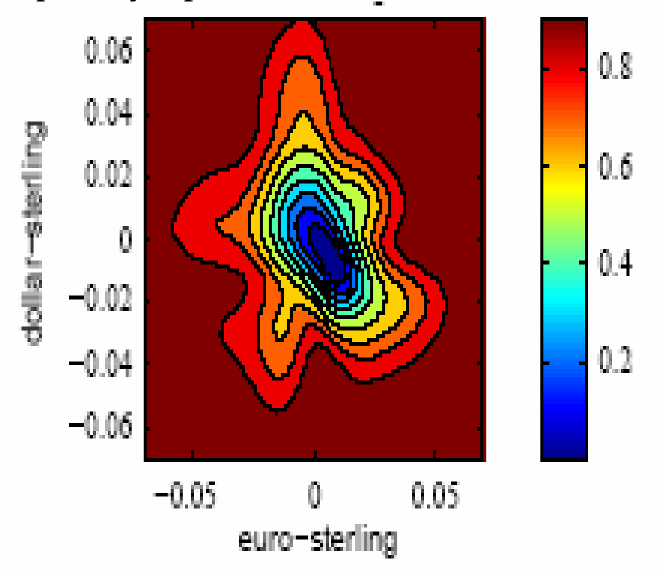

Chart 13b: Resulting unconditional distributions for (simplified) £ ERI.

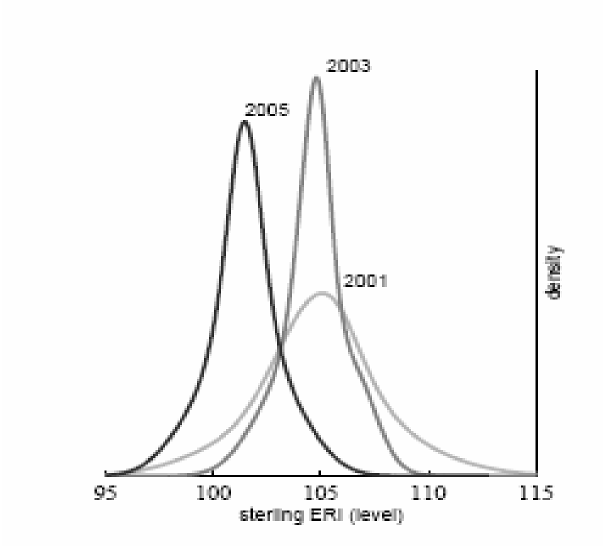

Accounting for the presence of uncertainty in derivatives

Having argued that (PDFs) derived from options prices provide a key tool for quantifying uncertainty and asymmetry attached to future levels of interest rates, asset prices and commodity prices. It is important to know how reliable a tool this is. As with other asset prices, risk premia provide a complication.

With options, the problem arises because the PDFs we derive reflect risk-neutral probabilities, whereas we are really interested in the underlying 'real world' probability distributions. This is potentially an important issue. We need to be confident that risk-neutral probabilities we use are a reasonable guide to agents' views Typically, these PDFs are used as if they represent the market's views on the distribution of the underlying asset, interest rate or commodity price. However, as the PDFs are extracted by assuming that investors do not require compensation for risk (i.e. they are 'risk neutral'), this is not the case. Risk premia, for example, mean that the central expectations differ.

Previous work has gone some way toward understanding the extent and nature of differences between these risk-neutral PDFs and probabilities actually held by market participants (see Bliss and Panigirtzoglou, 2004, for recent work at the Bank and

Taylor, 2005, for research elsewhere). Chart 14 below gives an example of how the presence of risk can alter the shape and location of the distribution in the circumstances when the underlying utility function is known. In practice, the implicit transformation or link function which transforms the risk neutral into the "real world" distribution is unknown and needs to be estimated[17]. We are currently examining this problem, including extending the basic empirical approach to include explicit time-varying adjustments for risk that are based on macro-variables and other observable factors.

Chart 14: Transformation of the risk-neutral density of the FTSE100 distribution (26/10/05) for a 3 month horizon (based on a CRRA utility function with coefficient gamma).

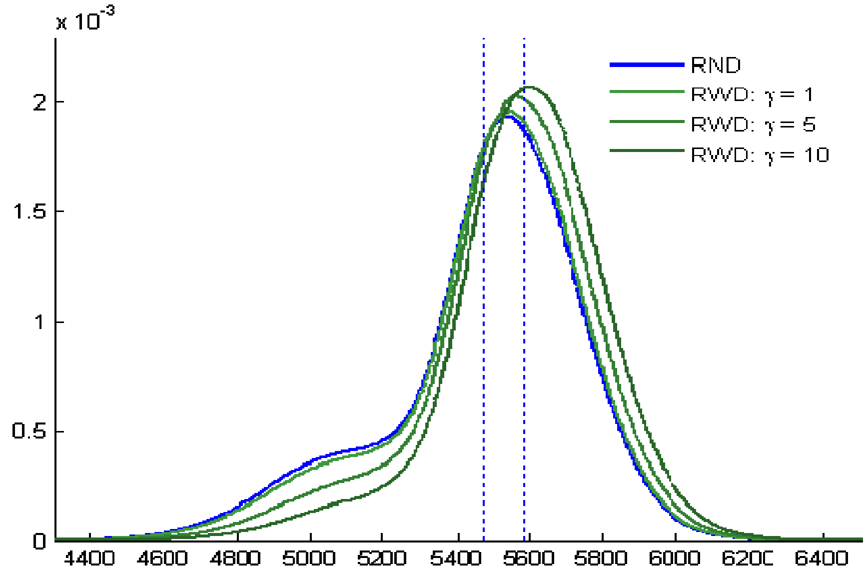

4. REFINING OUR UNDERSTANDING OF THE ROLE OF RISK

A common theme that has emerged from the description of the standard tools we use to extract information from asset prices is that risk premia play a central role. In particular, a stylised fact seems to be that asset prices move far more than can be justified by movements in the variables we think should determine them (ie by the "conventional macro fundamentals" such as income expectations or expected interest rates. To make matters worse, our understanding of how risk should affect asset prices

[17] An additional complication is that the risk neutral distribution may need to be adjusted to account for the presence of stochastic volatility or jump risk.

is rather sketchy. It is treated as a residual explanation rather than a fundamental driver of behaviour as it should be.

The last section has also described recent attempts to derive empirical estimates of risk premia. Considerable progress is being made in this area. But it is still the case that in many of the empirical models concerned, in particular those using the affine approach, the estimates derived are difficult to link to underlying economic factors. Some recent research has attempted to address this by explicitly incorporating macro factors in the empirical framework (eg Piazessi and Swanson, 2004, Rudebusch and Wu, 2003, and Lildholdt *et al* , 2006). But the fact remains that these attempts to link macroeconomics and finance are often partial and ad hoc.

An alternative approach to addressing this problem is to examine how risk and uncertainty impinge on macroeconomic behaviour and asset prices using a more rigorous theoretical approach. Currently, the general equilibrium macroeconomic models used in practice (for example BEQM at the Bank of England) does not explicitly incorporate consideration of risk and uncertainty.

So what needs to be added to a standard economic model in order to talk about risk and uncertainty like this? The answer is that no new *theory* needs to be added. All of this is implicit in a model like BEQM, for example. Instead, the difference lies in the solution methods that we are using. At the moment, it is only feasible to solve macro models of BEQM's size by imposing a property of *certainty equivalence*, in which agents act as if the future were known with certainty, once shocks have occurred. This has the effect of assuming risk away, or to be more accurate it subsumes the effects of risk premia in the constants of the model so as a result risk premia have no effects on the dynamic responses of the model.

But this simplifying assumption raises questions. What would macroeconomic projections look like if we were able to relax the certainty equivalence assumption? What would the world look like if uncertainty increased? Given the size of observed risk premia and their apparent tendency (as documented in section 3) to be time-varying, theory and practice suggests that we could be missing some important effects, but we do not have a good idea of what these might be or how large they could be?.

Fundamentals of risk premia

To understand how this general equilibrium modelling work is derived, it is important to remind ourselves of first principles. For this discussion, we will focus on the behaviour of imaginary consumer-investors, who are motivated to consume as much as possible over their lifetimes, and are able to buy assets to transfer wealth into the future. For simplicity, we assume here that these consumer-investors live forever and are not concerned with planning for different stages in their lives, such as retirement. (It also helps to assume that there are no market failures or frictions that distort asset prices.) We also assume, as is typical, that agents would like to smooth consumption, so that they are as consistently happy as possible.

Formally, we can express this in terms of a lifetime utility function for the representative consumer given by

$$U = E_t \sum_{i=0}^{\infty} \beta^i \frac{C_{t+i}^{1-\gamma} - 1}{1-\gamma},$$

where E is the expectations operator, C is consumption of a single perishable good produced by a single 'tree', $0<\beta<1$ is the subjective discount factor measuring the consumer's impatience and $\gamma>0$ is the coefficient of relative risk aversion.

Because these agents want to smooth their consumption, they care about the circumstances in which pay-offs are delivered. All things being equal, assets that are expected to pay well when growth and incomes are low will be more highly valued than assets that are expected to pay well in good states of the world. This latter type of asset is relatively *risky*, as its return is not earned when investors most need it; its return is paid out when the marginal utility from more consumption is low. Such an asset should sell at a discount to account for its riskiness, therefore driving up its expected return. This adjustment for riskiness is known as the risk premium, and it can be thought of as measuring the extra return an investor requires to be made indifferent between holding a riskless asset and a risky one.

Formally, for a general asset with return Ri, we can derive this result from the fundamental asset pricing equation (see Cochrane, 2004);

i.e $E_t[(M_{t+1} R_{t+1}^i)] = 1$,

where $M_{t+1} = \beta \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}}$ denotes the stochastic discount factor, from which it follows[18] that

$$E_t[R_{t+1}^i] = \frac{1}{E_t[m_{t+1}]} - \frac{\text{cov}_t[R_{t+1}^i, m_{t+1}]}{E_t[m_{t+1}]}$$

$$E_t\left[R_{t+1}^i\right] - R_{t+1}^f = -\frac{\text{cov}_t\left(m_{t+1,}R_{t+1}^i\right)}{E_t\left[m_{t+1,}\right]}$$

since the inverse of the expected SDF is the safe interest rate.

This expression for the risk premium is often usefully expressed as the product of the *price of risk* and *quantity of risk*, i.e.

$$E_t\left[R_{t+1}^i\right] - R_{t+1}^f = -\frac{cov_t\left(m_{t+1,}R_{t+1}^i\right)}{E_t\left[m_{t+1,}\right]} = -\frac{\rho_t^{m,R^i} \overbrace{\sigma_t\left[m_{t+1,}\right]}^{\text{volatility of SDF}} \sigma_t\left[R_{t+1}^i\right]}{E_t\left[m_{t+1,}\right]}$$

$$E_t\left[R_{t+1}^i\right] - R_{t+1}^f = -\underbrace{\frac{\text{cov}_t\left(m_{t+1,}R_{t+1}^i\right)}{\sigma_t^2\left[m_{t+1,}\right]}}_{\text{quantity of risk}\left(\beta_t^{i,m}\right)} \underbrace{\frac{\sigma_t^2\left[m_{t+1,}\right]}{E_t\left[m_{t+1,}\right]}}_{\text{price of risk}\left(\lambda_t^m\right)} = -\underbrace{\rho_t^{m,R^i} \frac{\sigma_t\left[R_{t+1}^i\right]}{\sigma_t\left[m_{t+1,}\right]}}_{\text{quantity of risk}\left(\beta_t^{i,m}\right)} \underbrace{\frac{\sigma_t^2\left[m_{t+1,}\right]}{E_t\left[m_{t+1,}\right]}}_{\text{price of risk}\left(\lambda_t^m\right)}$$

In a textbook model, the price of risk[19] is directly related to agents' attitudes towards risk. This may be summarised by the coefficient of relative risk aversion and the volatility of consumption growth. But where does the quantity of risk come from? The brief discussion above hinted that it is associated with the co-movement between an asset's return and economic conditions. More specifically, the quantity of risk is closely tied to the covariance between returns and the marginal utility an agent gains from extra consumption.

However, finance theory does not explain what determines the covariances of consumption and returns; the processes that generate pay-offs are usually taken as given. To understand risk premia more deeply, we need to understand what jointly determines consumption, incomes and asset pay-offs. This is the province of macroeconomics.

[18] Using the result that $E[xy] = E[x]E[y] + cov[x,y]$.

The interaction of macroeconomics and finance

In a dynamic macro model, the quantity of risk depends on the fundamental shocks and the structure of the economy. But simple textbook models produce risk premia that are too low, given realistic values of risk aversion. When we look at the covariance between returns and the growth in the marginal value of consumption (otherwise known as the consumption-based *stochastic discount factor*), we see it is not large enough. A first step towards realistic risk premia is to set up the model so that agents are very sensitive to changes in consumption. This is often done by assuming that agents get used to a certain level of consumption (ie they form a consumption habit), which increases their motivation to smooth consumption. Technically this is achieved by assuming a modified version of the utility function, of the form;

$$U = E_t \sum_{i=0}^{\infty} \beta^i \frac{(C_t - X_t)^{1-\gamma} - 1}{1-\gamma}$$

where X is the reference level of consumption. In the external habits model, this is given by aggregate consumption (see Campbell and Cochrane, 1999). This gives rise to a stochastic discount factor of the form;

$$M_{t+1} \equiv \beta \left(\frac{C_{t+1} - X_{t+1}}{C_t - X_t} \right)^{-\gamma} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \left(\frac{F_{t+1}}{F_t} \right)^{-\gamma},$$

where $F_t \equiv (C_t - X_t)/C_t$ is the 'consumption surplus ratio', so even if consumption is smooth, small shocks can generate variation in the SDF.

In finance theory, because incomes and asset returns are often assumed to be exogenous — as if fruit randomly falls off a tree, for example — these assumptions are able to generate realistic risk premia. But in macroeconomic models, incomes and asset returns are usually endogenous, and reflect how the supply side of the model economy reacts to fundamental shocks, such as to productivity. In this setting, it turns out that devices like consumption habits by themselves cannot generate realistic asset price behaviour. The reason is simple: if consumer-investors are able simply to eat into their stock of capital after a bad shock, and save in good times, they can effectively smooth their consumption as much as they want. In other words, it is not

[19] The inverse of the price of risk is sometimes known as "risk appetite"; see for example Gai and Vause, (2004)

enough that agents dislike consumption volatility; they have to be prevented from doing something about it. Unless we make agents incredibly risk-averse, such models will deliver low risk premia. This suggests that the behaviour of asset prices is telling us something important about the way goods and labour markets work. If goods and labour markets can adjust instantly after shocks, the economy will deal with uncertainty too easily, and risk premia will be unrealistically low. We therefore need frictions and rigidities in the real economy to make it harder for households to smooth consumption, thereby increasing the quantity of risk. In our modelling work we have confirmed the results by other authors (in particular Jermann 1998) who find that by introducing devices that prevent economic agents from smoothing consumption through altering production, empirically more plausible risk premia can be reintroduced. Adjustment costs are on such device that can create risk premia and suggest that friction in the model economies are likely be important if we want to make meaningful statements about risk and uncertainty in DSGE models.

Theory of term premia

We can apply the same logic used above to the term structure, which describes the rates of return on bonds of different maturities. In other words, term premia — the difference between the expected future interest rate and the forward rate for a given maturity — respond to the same macroeconomic factors as other risk premia, such as the equity risk premium. To see this, consider our consumer-investor again. When considering how to transfer wealth across two or more periods, the agent can either invest in consecutive one-period bonds, or invest in a long bond. The degree of risk in either strategy will depend on the evolution of returns and consumption over time. For example, if consumer-investors believe there is a probability that bad states of the world will be followed by more bad states, it makes sense for them to buy a long bond that ensures consumption in the future. In these circumstances, these bonds would serve as a type of insurance. This implies that term premia will be negative and the yield curve downward sloping. On the other hand, if consumer-investors believe the economy tends to revert strongly to a trend growth path, then there is less uncertainty associated with longer-term investments, and the yield curve will be upward sloping.

Monetary economies

So far, all our discussion has been based on real economies. In a monetary economy, the presence of inflation introduces another source of uncertainty. The premium agents require for being exposed to this is called the inflation risk premium. The inflation risk premium can be positive or negative. It depends on the co-movement between inflation and the real business cycle: in particular, if inflation co-varies negatively with consumption, the value of the inflation risk premium is *positive*. In this case, a nominal bond needs to pay a larger return, because high inflation reduces real returns when they are most needed.

Work on the implications of monetary policy and nominal rigidities in micro-founded dynamic macro models is currently in its infancy, and it is too early to draw wide-ranging conclusions. Nonetheless, the potential benefits include an improved understanding of the role of inflation risk premia in the yield curve, and, possibly, a new take on monetary policy rules; for example, how does the behaviour of the monetary authority contribute to macroeconomic risk?

Our initial findings in this area are described in De Paoli *et al.* (2006). The key results derived so far are as follows:

1. Risk premia are endogenous and reflect the way shocks on average work their way through the economy. So we should be looking at what risk premia tell us about fundamental macro shocks and the structure of the macroeconomy.
2. Equity and term premia are increasing in the degree of real rigidities (eg, capital adjustment costs). A higher degree of real rigidities means that the propagation of a given shock is relatively greater, meaning more volatility for marginal utility and asset returns.
3. Production frictions – rigidities that prevent rapid adjustment of capital and labour factors in response to shocks – are essential for achieving significant risk premia. Tricks used in finance (eg, consumption habits) are, by themselves, insufficient.
4. When there are only monetary policy shocks, the equity risk premium increases with the degree of nominal rigidity: nominal rigidities generate volatility in consumption and returns on capital. Both the nominal and real yield curves steepen with the degree of nominal rigidity.
5. When there are only productivity shocks, the equity risk premium decreases with the degree of nominal rigidity: nominal rigidities dampen the effects on quantities. Both curves flatten as the degree of nominal rigidity rises.

Chart 15: Stochastic mean of asset price indicators as degree of habit persistence is varied (taken from De Paoli et al, 2006).

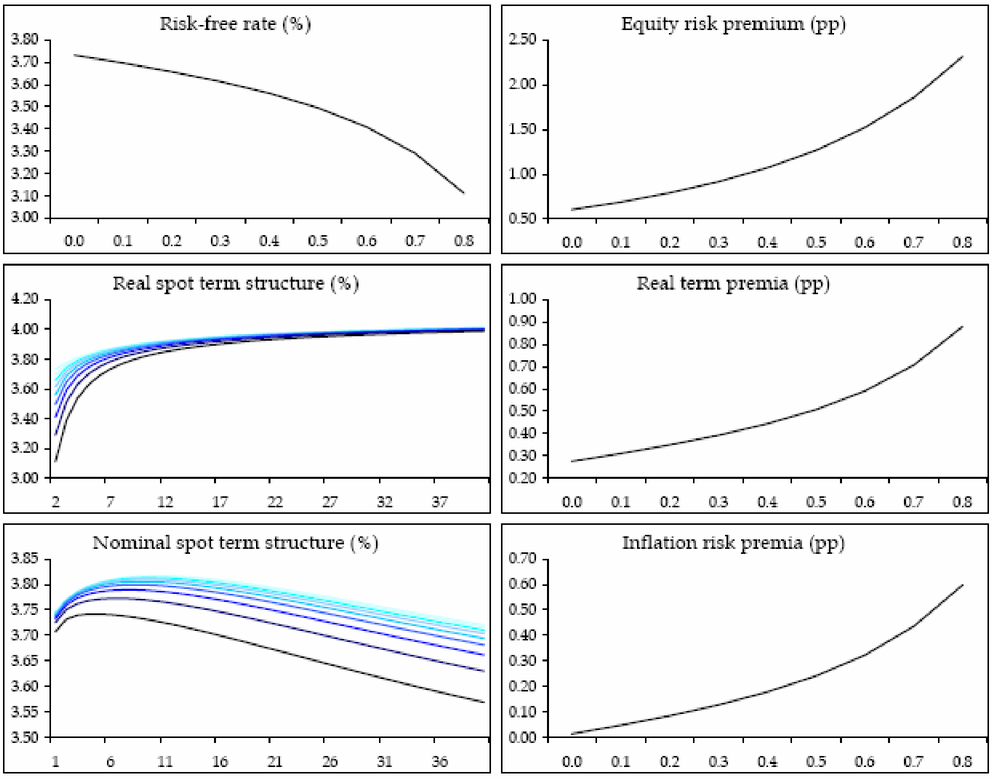

An application: The role of increased uncertainty in explaining the "Great Stability"

This theoretical work has also allowed us to derive important insights about current policy problems. During the period of inflation targeting in the United Kingdom, the volatility of output and inflation has been well below the levels seen in the previous four decades (see Table 2 below).

It has been argued that this so-called 'Great Stability' might also have contributed to a lower equity risk premium and lower real interest rates observed over this period.

Table 2 – Stylised facts business cycle and asset price facts

	Standard deviation		*Mean*					
	Real output growth[(a)]	Inflation[(a)]	Real risk-free rate[(b)]	Real 10-year spot yield[(c)]	Real 10-year spot yield[(d)]	Real stock return[(e)]	Equity risk premium[(f)]	Yield spread[(g)]
1960-69	2.0	1.6	2.0	2.6	n/a	8.6	6.6	0.6
1970-79	2.5	5.7	-3.6	-0.6	n/a	4.4	8.0	3.0
1980-92	2.5	3.9	4.1	4.3	3.9	13.0	8.8	0.1
1993-04	0.8	0.4	2.9	3.2	2.7	7.6	4.7	0.3

(a) Standard deviation of annual percentage change, (b) end-quarter annual return of three-months T-bill minus end-quarter annual inflation rate, (c) end-quarter 10-year benchmark nominal spot yield minus proxy for 10-year inflation expectations (centred 10-year moving average of end-quarter annual inflation rate), (d) end-quarter VRP 10-year real yield, (e) end-quarter annual FTSE All-Share Total Return minus end-quarter annual inflation rate, (f) difference between data described under (e) and (b), (g) difference between data described under (c) and (b).

In fact, using an earlier simplified version of the model described in De Paoli et al., (2006) our theoretical analysis of the role of uncertainty in affecting risk premia and other macroeconomic variables led us to conclude that this conclusion may be too simplistic:

- In our modelling framework, a lower volatility of shocks indeed lowers the volatility of real output growth, the equity risk premium and term premia. This is consistent with the stylised facts. Chart 16a and b illustrates for a range of assumed values for the volatility of productivity shocks.

Chart 16a[(a)]

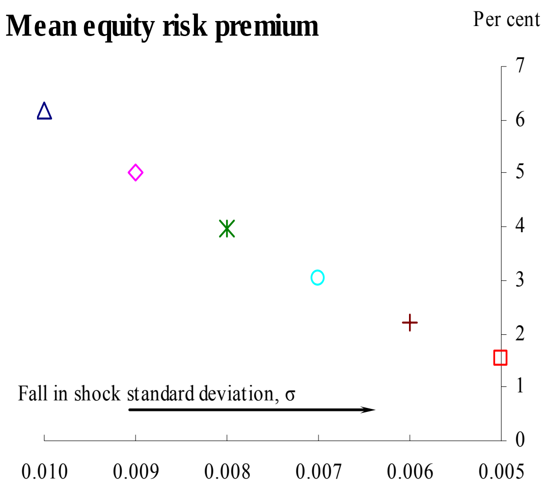

Chart 16b[(b)]

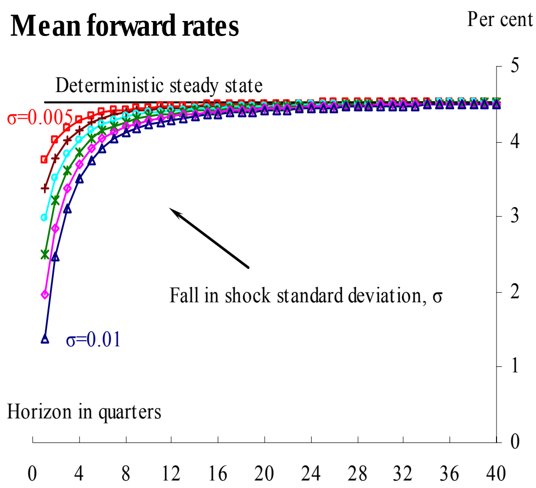

(a) The mean equity risk premium is the difference between the annualised mean of the return on a one-quarter real bond (the risk-free rate) and the annualised mean of the return on equity, (b) calculated from mean of quarterly spot yields.

But by reducing the precautionary savings motive, a lower volatility of shocks also increases the real risk-free rate. This rise outweighs lower term premia, with the overall effect of lower volatility being *higher* forward rates at all horizons. This is inconsistent with the stylised facts (this point has also been made by Ferguson, 2006). Chart 17a and b illustrates. So for the type of model considered here, the effect on the real risk-free rate and expected future real risk-free rates will always dominate the effect on forward premia.

Chart 17a

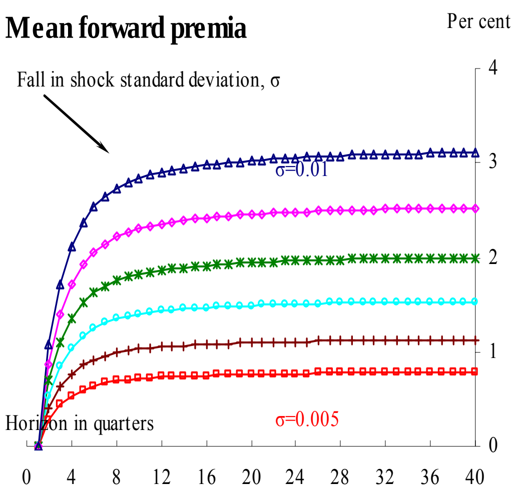

Chart 17b

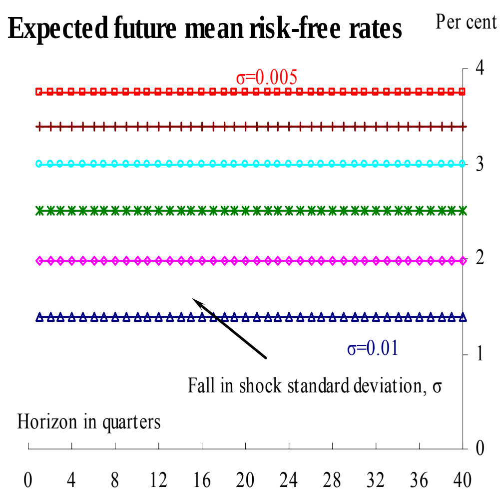

- Although the simple real business cycle model is silent about the effects of better fiscal or monetary policy, a reduction in the volatility of fiscal and/or monetary shocks alone would likely produce similar qualitative results. But risk premia also reflect the structure of the economy. It is possible that the fall in the volatility of the shocks reflects structural changes and by extension risk premia. This might reconcile lower output volatility and lower real interest rates so further investigation of this channel is necessary.

5: UNDERSTANDING ASSET PRICE MOVEMENTS WHEN THE TEXTBOOK MODEL FAILS

The last section has described our attempts to rationalise risk premium movements in terms of macroeconomic fundamentals. But inevitably, in practice, it is still likely to be the case that asset price movements are not possible to rationalise in terms of "conventional fundamentals" or in terms of plausible developments in risk premia.

In such circumstances, it is important to know whether such asset price movements are distorted by market-related factors not associated with the fundamentals of the textbook model. This can occur because markets are segmented, or at least imperfectly integrated. So technically, there may be imperfect substitutability across assets and markets. This may be because there are not enough investors who have sufficiently long holding periods to bring forward prices into line with expectations. This argument is often used when talking about demand-supply imbalances caused by pension fund asset-liability management or Asian central bank purchases driving down yields.

Illustration of "non-textbook" effects on asset prices

There are numerous examples of episodes where potential demand-supply have appeared to influence asset prices. Here, we consider the experience of UK inflation breakevens in late 90s.

Chart 18: UK implied 12-month forward inflation rates

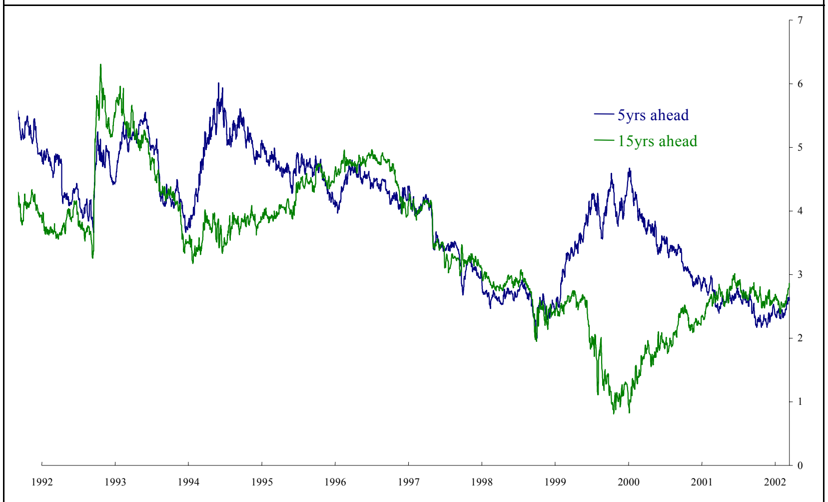

On the day the UK Government announced operational independence in monetary policy for the Bank of England these implied future inflation rates fell by up to about half a percentage point. They then drifted down to around 2½ % over the next year. But through 1999 implied inflation rates for a few years ahead increased again to over 4% even though there was no such increase in direct surveys of inflation expectations. And bond-implied inflation rates for the more distant future fell to remarkably low levels (sometimes as low as 1 %). Subsequently, future inflation rates implied by relative bond prices have come closer together again at around 2½%.

What can explain these movements? Institutional investors including pension funds are large holders of British Government bonds. The Minimum Funding Requirement (MFR) for pension funds was coming into effect in the late 1990s and then possible reform to the MFR was increasingly discussed. It seems quite likely that these developments had a significant impact on bond yields so that the apparent economic information from the market was distorted by institutional factors.[20]

7: CONCLUSIONS

This paper has argued that asset prices provide a rich source of informative and timely information of relevance to central bank policymakers. A range of modelling tools has been presented which has been designed to allow information about the real shocks hitting the economy to be reverse-engineered from asset prices.

A central challenge in the interpretation of asset prices is the identification and interpretation of risk premia. Recent research at the Bank of England has been attempting to decompose asset price movements into those associated with standard macroeconomic fundamentals and those caused by shifting perceptions or attitudes towards risk. It has been argued that it is important to take a dual approach to this problem whereby the empirical estimates of risk premia are complemented by a sound theoretical understanding of how risk premia might be determined in a rigorous framework which explicitly recognises the interactions between macroeconomics and finance.

[20] See Cooper and Scholtes (2001).

REFERENCES

Anderson , N and Sleath, (1999 and 2001), 'New estimates of the UK real and nominal yield curves', Bank of England Quarterly Bulletin, November 1999, pages 384-92, and Bank of England Working Paper, No 126.

Backus, D,K, Foresi, S, and Telmer,C,I, `Affine Term Structure Models and the Forward Premium Anomaly', The Journal of Finance, 56, No 1, pp 279-304.

Bekaert G, Engstrom E and Xing Y (2005), 'Risk, uncertainty and asset prices', Federal Reserve Board, Finance and Economics Discussion Series

Benati, L. (2006) 'Affine term structure models for the foreign exchange risk premium', Bank of England Working Paper, No. 291.

Bernanke, B. (2004) 'What Policymakers Can Learn from Asset Prices' Remarks before The Investment Analysts Society of Chicago, Illinois, April 15.

Bliss, R. and N. Panigirtzoglou, 2004, Recovering risk aversion from options, Journal of Finance, 59, 381-422.

Breedon D & Litzenberger R, Prices of state-contingent claims implicit in option prices. Journal of Business, 51, pp621-651

Brooke, M, Cooper, N and Scholtes, C (2000), 'Inferring market interest rate expectations from money market rates', Bank of England Quarterly Bulletin, November, pages 392-402.

Brigden, A, Martin, B and Salmon, C (1997), 'Decomposing exchange rate movements according to the uncovered interest parity condition', Bank of England Quarterly Bulletin, Vol. 37 No. 4,pages 377-88.

Campbell J Y and Cochrane, J H (1999), 'By force of habit: a consumption-based explanation of aggregate stock market behavior', Journal of Political Economy, Vol. 107, No. 2, pages 463-506

Clare, A and Lekkos, I (2000), 'An analysis of the relationship between international bond markets', Bank of England Working Paper, No 123.

Clews, R. (2002) 'Asset prices and inflation', Bank of England Quarterly Bulletin, Summer.

Clews, R, Panigirtzoglou, N and Proudman, J (2000), 'Recent developments in extracting information from options markets', Bank of England Quarterly Bulletin, February, pages 50-60.

Cochrane, J.H. (2004), Asset Pricing, Princeton Unviersity Press, Princeton NJ, Revised Edition.

Cochrane, J (2006), 'Financial markets and the real economy', manuscript, http://faculty.chicagogsb.edu/john.cochrane/research/Papers/index.htm

Cooper, N and Scholtes, C (2001), 'Government bond markets in an era of dwindling supply' in 'The changing shape of fixed income markets: a collection of papers by central bank economists', BIS Paper, No 5.

De Paoli, B., Scott, A. and Weeken, O. (2006), 'Asset pricing implications of a New Keynesian model', Bank Working Paper forthcoming.

Ferguson, R.W. (2005) Remarks to the Banco de Mexico International Conference, Mexico City, Mexico, November 15; http://www.federalreserve.gov/boarddocs/speeches/2005/200511152/default.htm

Gai, P and Vause, N (2004), 'Risk appetite: concept and measurement', Bank of England Financial Stability Review, December, pp. 127-136.

Harrison, R., Nikolov, K, Quinn, M., Ramsay, G., Scott, A. and Thomas, R. (2005), The Bank of England Quarterly Model, Bank of England.
Hurd, M. and Relleen, J. (2006), 'New information from inflation swaps and index-linked bonds', Bank of England Quarterly Bulletin, Vol. 46, No. 1.
Hurd, M., Salmon, M. and Schleicher, C. (2006) 'Using Copulas to Construct Bivariate Foreign Exchange Distributions with an Application to the Sterling Exchange Rate Index', Bank of England Working Paper forthcoming.
Jermann, U (1998), 'Asset pricing in production economies', Journal of Monetary Economics, Vol. 41, pages 257–75.
Kim, D.H. and Wright, J.H. (2005), 'An arbitrage-free three factor term structure model and the recent behaviour of long-term yields and distant horizon forward rates', Federal Reserve Board Working Paper No. 2005-33.
Lildholdt, P., Peacock, C. and Panigirtzoglou, N. (2006), 'An affine macro-factor model of the UK yield curve', Bank of England Working Paper forthcoming.
Norges Bank (2005), Inflation Report 3/2005 .
Panigirtzoglou, N and Scammell, R (2002), 'Analysts' earnings forecasts and equity valuations', Bank of England Quarterly Bulletin, Spring, pages 59-66
Piazessi, Monika, and Eric Swanson (2004). 'Futures Prices as Risk-Adjusted Forecasts of Monetary Policy,' University of Chicago and Board of Governors of the Federal Reserve System, Working paper (March).
Svensson, L. 'The Inflation Forecast and the Loss Function', in Paul Mizen, ed. (2003), Central Banking, Monetary Theory and Practice: Essays in Honour of Charles Goodhart, Volume I, Edward Elgar, 135-152.
Taylor, S.J., 2005, Asset Price Dynamics, Volatility, and Prediction, Princeton University Press.
Wells S (2005), 'Real interest rates and macroeconomic volatility', Bank of England Quarterly Bulletin, Summer.