Deriving a market-based measure of interest rate expectations

By Christopher Peacock of the Bank's Monetary Instruments and Markets Division.

Forward rates are perhaps the most common measure of expected future interest rates. But the existence of a risk premium can drive a wedge between forward rates and what the market expects future rates to be. In this article we use survey data to derive an estimate of the risk premium. We find that the survey-based risk premium implies a significant and time-varying difference between forward rates and expected future interest rates. Consequently, this article sets out a simple model of the survey-based risk premium that can be used to generate a path for expected future interest rates on any particular day.

Introduction

In a speech two years ago, the then Governor urged: '*... if I could make just one comment on the current macroeconomic situation, ... it would be to caution you against placing too much weight on the steepness of the short-term interbank [forward] curve as an indicator of the likely course of official short-term interest rates ...*'.(1) Chart 1 plots the forward curve on the eve of the Governor's speech.(2)(3) It shows the interbank forward rate rising from a little over 4% to around 5.4% a year later. But in a poll of financial market economists, published just one week earlier, the average expectation of the interbank rate a year ahead was 5.0%. In other words, there was a difference of some 0.4 percentage points between the market forward rate and what these economists, on average, expected the interbank rate to be.

This article begins by discussing why the presence of a risk premium can lead to forward rates being a biased measure of expected future interest rates. The third section shows how surveys of interest rate expectations can be used to derive an estimate of the risk premium. It finds that the survey-based risk premium implies a significant and time-varying difference between forward rates and expected interest rates. In light of this, the fourth section proposes a simple model of the survey-based risk premium that can be used to generate a path for future interest rate expectations on any particular day. The fifth section applies this model to examine what the estimated profile for interest rate expectations would have been at 18 February 2002, and also at 28 May 2004.(4) The last section concludes.

Chart 1
Short-term interbank forward curve at 18 February 2002

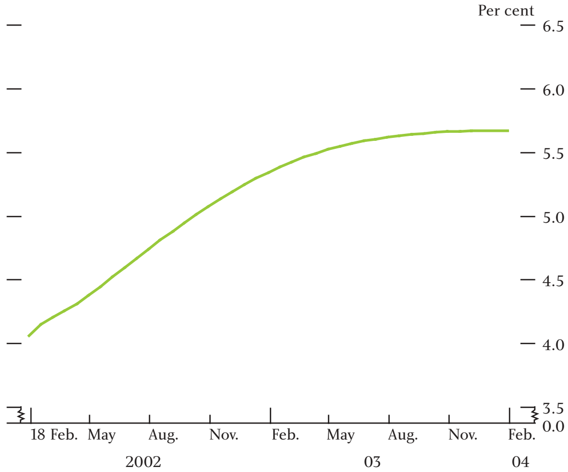

Forward rates and expectations of future interest rates

The premise that the forward curve represents the path of expected future interest rates is known as the *expectations hypothesis*. But in practice there are a number of factors that may drive a wedge between forward rates and what the market expects future rates to be. For instance, if market participants are risk averse

(1) See George (2002).
(2) These forward rates are available daily at www.bankofengland.co.uk/statistics/yieldcurve/main.htm.
(3) Forward rates are the interest rates for future periods that are implicitly incorporated within today's interest rates for loans of different maturities. For instance, suppose that the interest rate today for borrowing for six months is 6% per annum and that the rate for borrowing for twelve months is 7% per annum. Combined, these two interest rates contain an implicit interest rate for borrowing for a six-month period starting in six months' time of roughly 8% per annum.
(4) The data cut-off for the 'Markets and operations' article in this *Quarterly Bulletin*.

they are likely to require a *term premium* as compensation for the uncertainty about future interest rates. Participants may also demand a *liquidity premium* to hold instruments that are difficult to trade at times of market stress. Both these factors, which we collectively refer to as risk premia, will tend to push the forward curve above the path of expected future interest rates.

Ideally, one would test the expectations hypothesis by comparing forward rates directly with expectations. However, expectations are unobservable. To get round this problem, empirical studies often assume that expectations are *rational*.(1) In other words, that expectations of future interest rates do not differ systematically from subsequent interest rate outturns. By assuming rational expectations hold, any systematic difference between forward rates and outturns can be interpreted as reflecting the risk premium.

Chart 2 shows the differences between interbank forward rates and corresponding outturns for a range of horizons out to two years, over the period May 1993 to April 2004. Each bar shows the range of the differences at a particular horizon, with the average shown by a red square. As noted above, by assuming rational expectations, the expectations hypothesis suggests that these differences should average zero. By contrast, the chart shows that, on average, interbank forward rates have provided an upwardly biased forecast of future interbank rates. Over the period, the average differences at three months, one year and two years ahead were 14, 66 and 135 basis points respectively. The chart also makes clear the often large differences, both positive and negative, between interbank forward rates and subsequent outturns. Moreover, the range of these differences increases with the horizon, probably reflecting the greater level of uncertainty as market participants project further out into the future.

On this basis at least, it would appear that the expectations hypothesis can be rejected for the interbank market. But the above approach has one clear drawback. It assumes that expectations are rational. However, since expectations are unobservable, there is no way to test whether this assumption is valid. This means that the biases observed in Chart 2 may be due not only to the existence of a risk premium, but also to market participants making systematic expectational errors.(2)

Chart 2
Differences between interbank forward rates and subsequent outturns

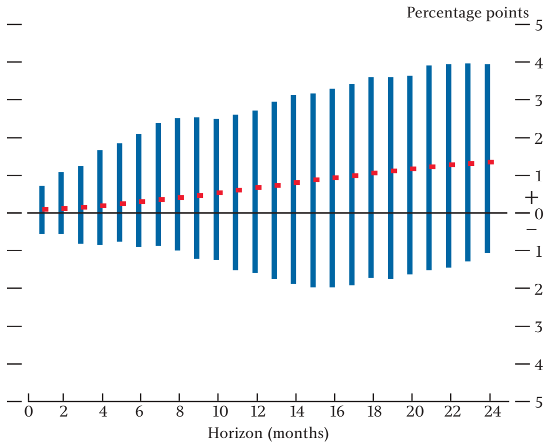

Survey expectations

Other empirical studies have used surveys of expectations as a proxy for the market's true expectations.(3) It can be argued that there are a number of advantages to using survey data as a means of estimating the risk premium. First, *ex-post* measures, such as the one shown in Chart 2, only provide an estimate of the *average* risk premium at a particular horizon. By contrast, surveys provide a time-varying estimate. And, as we show later on, there is a good deal of evidence to suggest that the risk premium does vary significantly over time. Second, survey-based estimates of the risk premium are immune to the impact of shocks that might occur between the survey date and the outturn. Thus, unlike the differences plotted in Chart 2, they are not affected by expectational errors. Third, by using survey data one can test for, rather than simply assume, rational expectations.

In this article we use surveys of short-term interbank rate expectations conducted by *Consensus Economics*. On a monthly basis it polls around 20 financial market economists on their expectations of the end-month interbank rate, both for three months and one year ahead. For instance, in the survey published on 8 April 2004, forecasts were reported for end-July 2004 and end-April 2005.

In terms of the survey data, rational expectations mean that survey respondents should not make systematic errors when formulating their expectations. The most

(1) See, among others, Fama and Bliss (1987) and Campbell and Shiller (1991).
(2) Another drawback is that, even if expectations are rational, it may be that the sample period is too short for non-systematic expectational errors to average out to zero.
(3) See Froot (1989) and MacDonald and Macmillan (1994). For a survey of the literature see Maddala (1991).

obvious way respondents might make systematic errors is if, over a long enough time period, their forecast errors are biased. Chart 3 plots each set of survey expectations data against its corresponding outturn. If the surveys are unbiased, the scatter plot should be centred on a 45-degree line that passes through the origin. The chart shows that there is no apparent bias in the survey expectations at either horizon, a finding confirmed by the results of a formal test outlined in the Appendix.

Chart 3
Survey expectations and outturns

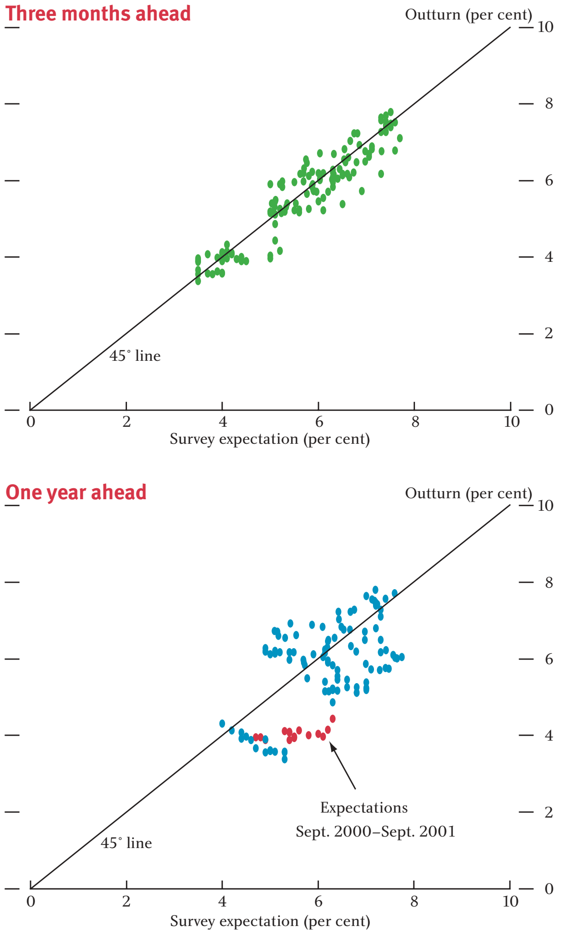

Nevertheless, it is clear from the chart that, especially at the one-year horizon, there are a number of points where interbank outturns were considerably below survey expectations. But this finding may be misleading due to the *overlapping observations problem*. Overlapping observations occur when the time between survey dates is less than the forecast horizon. The result is that a shock to the variable being forecast will affect the forecast errors over several periods, and not just one. For instance, the terrorist attacks on the United States in September 2001 led to sharp falls in equity prices, and a decline in business and consumer confidence. In response, policymakers both at home and abroad reduced official policy rates. Though the impact of this unforeseen shock affected short-term interest rates in only one month, it meant that the expectations for one year ahead taken between the Septembers of 2000 and 2001 would, other things being equal, be too high. Indeed, this was the case, as the expectations plotted in red testify.

A second reason why survey respondents may make systematic errors is if they fail to use all the available information at hand when formulating their expectations. If this were the case, then their expectational errors may be systematically related to the information they ignored. In the Appendix we test this *efficiency* condition by examining whether expectational errors were related to the level and the slope of the forward curve at the time the survey expectations were set. We use information from the yield curve since it can be viewed as a summary measure of potentially relevant explanatory factors, such as expectations of future inflation and output growth. The results of the test suggest that survey expectations were indeed efficient with respect to this information. So, in conjunction with the finding that survey expectations are unbiased, there is no strong evidence to suggest that the survey expectations we use in this article were not rational.

From the survey data we can derive an estimate of the risk premium as simply the difference between the interbank forward rate and the survey expectation of the interbank rate. At this stage it is worth noting how different this measure of the risk premium can be compared with the measure shown in Chart 2, and thus one that assumes expectational errors were zero at all points in the past. By way of an example, Chart 4 shows time profiles for the survey-based risk premium and the difference between the interbank forward rate and subsequent outturn, both at the one-year horizon. It is clear from the chart that there are a number of major differences between the two measures. First, over most of the period shown, the variation of the survey-based risk premium across time is considerably less than the *ex-post* difference. Second, the two measures do not track each other particularly well. Looking over the sample as a whole, the correlation between the two series is negative and around 0.2. Third, the two measures can lead to very different conclusions about the size of the risk premium. For instance, since the beginning of 2000 the *ex-post* difference has averaged over 90 basis points, compared with just 5 basis points for the survey-based risk premium.

Chart 4
Survey-based risk premium and difference between interbank forward rate and subsequent outturn: one year ahead

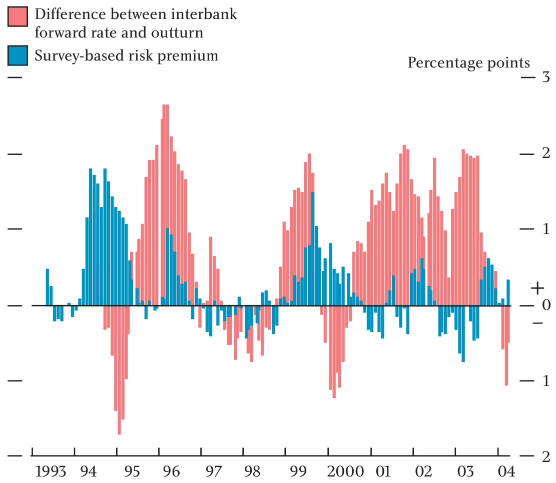

Chart 5
Survey-based risk premium

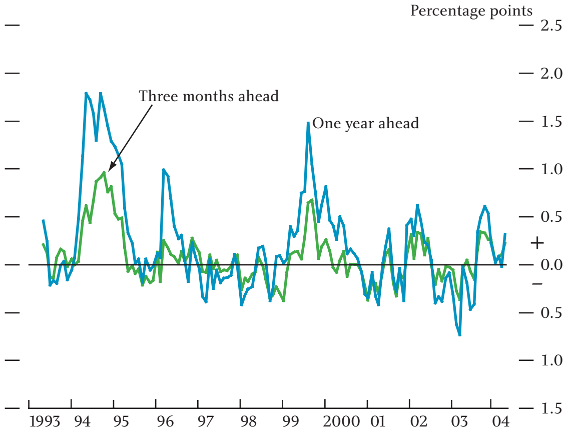

The analysis above is somewhat informal, but it does illustrate the inherent problem in using *ex-post* measures of the risk premium. Put simply, expectational errors can often swamp the information contained in *ex-post* measures, even at relatively short horizons. Consequently, in this article we focus on deriving, and then modelling, a survey-based, *ex-ante* measure of the risk premium.

By assuming rational expectations, we found evidence that appeared to suggest that the expectations hypothesis does not hold in the interbank market. What conclusion do we draw if we use survey expectations? In the Appendix we present the results of a formal test of the expectations hypothesis. In short, our conclusion is the same—the expectations hypothesis appears not to hold. To provide some intuition for why this is the case, Chart 5 plots the survey-based risk premium at three months and one year ahead. At both horizons the estimated risk premium is, on average, positive, at 7 and 22 basis points respectively. More importantly, the survey-based risk premium displays considerable variation over time, with peaks over eight times the sample average. In particular, at both horizons the estimated risk premium widened markedly in the middle of 1994 and 1999. Perhaps surprisingly, the chart shows that just under a half of estimated risk premia outturns are negative. It may be that our survey expectations provide, on average, an overestimate of the market's true expectation. But, as we discuss later on, there may also be a theoretical basis for the negative risk premia we observe.

Modelling survey-based risk premia

The survey data we employ in this article provide us with a *monthly* estimate of the risk premium. But with fast-moving financial markets such estimates can become quickly out of date. Consequently, the aim of this section is to model the time variation in the risk premium using variables that are available to us on a *daily* basis. In turn, this will enable us to provide an estimate of the risk premium and, therefore, interest rate expectations on any particular day.

Slope of the yield curve

A number of studies have documented a close link between the slope of the yield curve and measures of the risk premium.(1) Empirically, the risk premium is found to be positive and high when the yield curve is steep, and low when the yield curve is flat. Moreover, measures of the risk premium are often observed to be negative when the yield curve is downward sloping.

One explanation for this relationship comes from the literature on habit formation.(2) In this literature, agents' risk preferences are affected by the economic cycle through its impact on aggregate consumption relative to some habit level. In particular, habit formation models suggest that at the bottom of the cycle, when consumption is relatively low, risk aversion and risk premia tend to be high. At the same time, the yield curve tends to be upward sloping in anticipation of future rises in short-term interest rates. At the top of the

(1) See, among others, Fama and French (1989).
(2) See Campbell and Cochrane (1999) and Wachter (2004).

cycle the reverse is true. When consumption is relatively high, risk aversion and risk premia are likely to be low. At the same time, the yield curve tends to slope downwards.

Another explanation comes from the impact that expected interest rate changes may have on investors' appetite for taking on *interest rate risk*. This is the risk that movements in the yield curve may lead to capital losses. Specifically, if, relative to their central expectation, investors place a larger probability on rises in future interest rates than on falls in interest rates when the yield curve is upward sloping, then they are likely to require a positive risk premium to compensate them for the greater risk of future capital losses.

Chart 6 plots the survey-based risk premium one year ahead together with the slope of the yield curve. The chart does imply a fairly close relationship. In particular, when the survey-based risk premium was negative, the yield curve also tended to be downward sloping.

Chart 6
Survey-based risk premium and slope of the yield curve(a)

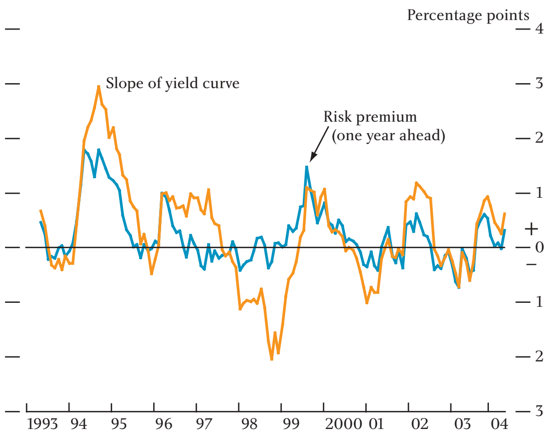

(a) Calculated as the difference between the two-week end-of-month government forward rate one year ahead and the official repo rate.

Interest rate uncertainty

It is likely that the more uncertain investors are about future asset returns, the greater the compensation they will require to hold risky assets. Indeed, this trade-off between risk and return is one of the key foundations of modern financial economics. Fortunately, the prices of some financial instruments, such as options, imply an expectation of the risks around future asset returns.

Consequently, Chart 7 plots a measure of expected interest rate volatility three months ahead, derived from options prices. There does appear to be a close empirical link between the survey-based risk premium and expected volatility. In particular, the peaks in expected volatility seen in the middle of 1994 and at the end of 1999 were matched by relatively high levels of the survey-based risk premium.

Chart 7
Survey-based risk premium and expected volatility

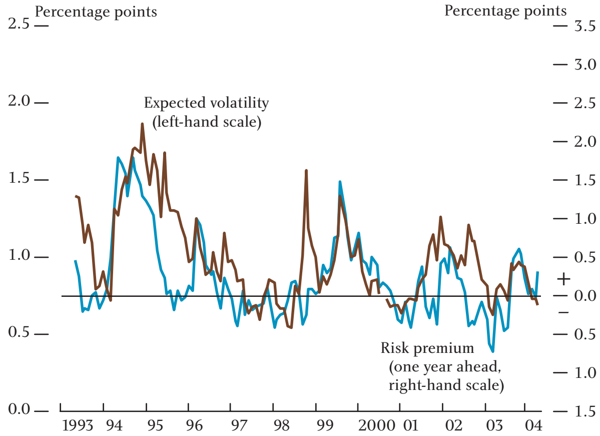

Measures of liquidity premia

Investors are often willing to accept a lower yield on assets that are more liquid and, therefore, easier to trade at times of market stress. Consequently, the liquidity premia attached to illiquid assets, such as interbank deposits, may serve to push forward rates above expectations of future interest rates.

Empirical studies often measure the liquidity premium as simply the difference in yield between two assets that have different liquidity, but are otherwise closely matched in terms of maturity, cash flow and credit risk.(1) These studies find that liquidity premia often vary considerably over time, widening markedly at times of extreme market stress. Such episodes are commonly termed *flights to liquidity*, for example the developments in Autumn 1998 in response to the Russian debt default.

In our model, we include two measures of liquidity premia to account for possible distortions across the forward curve. To capture the impact of changes in liquidity premia on the underlying interbank rate, and thus on the short end of the forward curve, we use the

(1) Examples include the spread between Treasury bonds and Treasury bills (Kamara (1994)), and *on-the-run* and *off-the-run* government bonds (Krishnamurthy (2002)). On-the-run bonds are the most recently issued bonds of a particular maturity. As these bonds are more frequently traded than off-the-run bonds, they are typically more expensive and therefore carry a slightly lower yield.

difference in yield between three-month certificates of deposit (CDs) and interbank deposits. These two money market instruments are subject to the same credit risk, but, unlike interbank deposits, CDs are traded in secondary markets and so enjoy a small liquidity premium.

To account for the impact that flights to liquidity may have on the longer end of the forward curve we use a five-year swap spread.(1) Empirically, changes in liquidity premia are often found to be an important factor in explaining the path of swap spreads.(2) This is borne out by Chart 8. It shows that the five-year swap spread widened markedly in the wake of the Russian debt default and subsequent collapse of the Long Term Capital Management (LTCM) hedge fund in 1998.(3)

Chart 8
Survey-based risk premium and swap spread

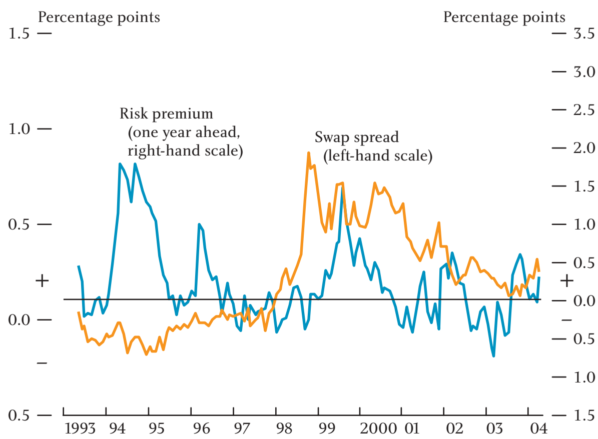

Following the discussion above, our model regresses the survey-based risk premium on the following four variables (together with a constant):

- the slope of the yield curve (denoted *Slope*);
- the expected volatility in the interbank rate three months ahead (denoted *Vol*);
- the spread between interbank deposit and three-month certificate of deposit rates (denoted *CD*); and
- the five-year swap spread (denoted *Swap*).

It is possible that use of the variables above will lead to an *endogeneity bias* in the model. For example, a positive shock to the risk premium will also tend to push up the slope of the yield curve. If this is the case, the slope of the yield curve will be positively correlated with the equation error, and the estimated regression coefficients will be biased. To account for possible endogeneity bias we use an estimation method known as *instrumental variables*. This method uses variables (called *instruments*) that are correlated with the endogenous variable, but are predetermined and thus uncorrelated with the equation error. This ensures that the regression coefficients are estimated consistently. In our model the instruments chosen are the regression variables lagged by one day.

Table A shows the results of the regression over the period May 1993 to April 2004, with standard errors in brackets, and significant parameters in bold.(4)(5) The model appears to provide a close fit, with the explanatory variables capturing over 70% of the variation in the survey-based risk premium, as measured by the regression correlation coefficient (denoted R^2). All variables have the expected sign, and the slope and swap measures are significant at both horizons. By contrast, expected volatility is found not to be an important factor at the three-month horizon, while the CD spread is not found to be important at the longer horizon. Note that all variables are measured in percentage points. This means that a 1 percentage point increase in the slope of the yield curve will lead to an estimated increase in the survey-based risk premium of 0.26 percentage points at the three-month horizon, and 0.42 percentage points at the one-year horizon.

Table A
Modelling the survey-based risk premium

	Three months	Twelve months
Const.	-0.10 (0.08)	**-0.52** **(0.20)**
Slope	**0.26** **(0.03)**	**0.42** **(0.06)**
Vol	0.03 (0.09)	**0.44** **(0.21)**
CD	**1.01** **(0.52)**	1.70 (1.17)
Swap	**0.16** **(0.08)**	**0.57** **(0.14)**
R^2	0.74	0.72

(1) A swap spread is the difference between a swap rate and a government bond yield of the same maturity.
(2) See Duffie and Singleton (1997) and Liu *et al* (2002).
(3) For a more detailed discussion of swap spreads and the factors that drive them, see Cortes (2003).
(4) Standard errors are calculated using Hansen's (1982) generalised method of moments to correct for overlapping observations.
(5) Bold variables are significant at the 10% level.

Chart 9 shows the contribution of each explanatory variable to the survey-based risk premium at the one-year horizon. The residual indicates the extent to which the model fails to capture the survey-based risk premium exactly. It is clear that no single explanatory variable has dominated movements in the fitted risk premium. Nevertheless, Chart 9 shows that, across our two measures of liquidity premia, the swap spread has played an increasing role in the second half of the period. By contrast, the influence of the CD spread has diminished markedly.

Chart 9
Contributions to the survey-based risk premium one year ahead(a)

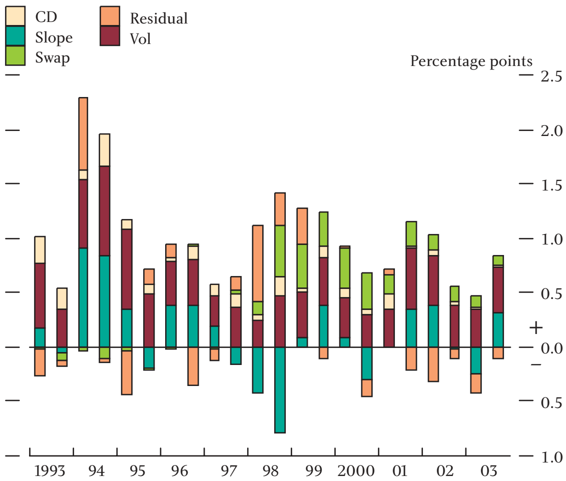

(a) For ease of exposition only the contributions for June and December in each year are shown. The constant is also not shown.

Simple regression models of the type outlined above are by their nature susceptible to instability resulting from changes in the underlying economic structure. For instance, it may be that the relationship outlined in our model is sensitive to changes in the level of the interest rate, or perhaps to the underlying monetary regime. To shed some light on the stability of the model, Chart 10 plots recursive estimates and standard error bands for the slope of the yield curve parameter.(1) The chart shows that, after an initial period of variability, the estimated relationship between the slope of the yield curve and the survey-based risk premium has been broadly stable since 1997.

One drawback of the model is that it allows us to adjust the forward curve only at two points, namely at three months and one year ahead. Fortunately, on a quarterly basis, Consensus Economics surveys expectations of the end-quarter interbank rate up to seven quarters ahead. For example, the quarterly survey published on 8 March 2004 reported end-quarter forecasts for 2004 Q1 through 2005 Q4.

Chart 10
Recursive estimates of the slope of the yield curve parameter in the model

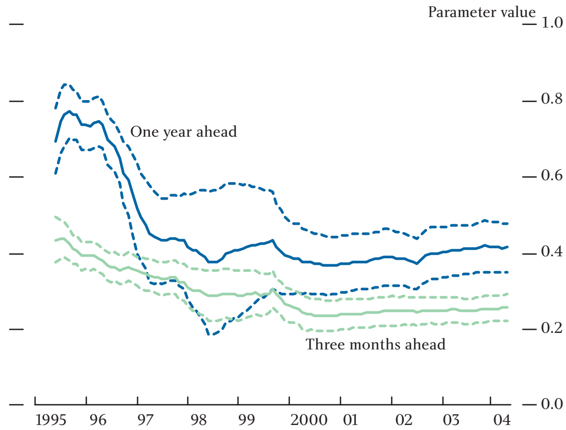

The Appendix reports the estimation results from regressing the survey-based quarterly risk premium on the explanatory variables. In short, the results are qualitatively unchanged: the quarterly model provides a good fit of the survey-based risk premium, and the explanatory variables are generally found to be significant and of a similar magnitude to those in the monthly model.

Adjusting forward curves for risk premia

By combining the estimated parameters from the monthly and quarterly models we are able to generate a path for future interest rate expectations on any particular day. By way of an example, Chart 11 shows the market forward curve plotted in Chart 1 for 18 February 2002 together with the adjusted forward rates joined by a cubic spline.(2) The adjusted forward curve suggests a much shallower path for expectations of future interbank rates than that embodied in the market forward curve. In particular, it implies that the interbank rate was expected to be around 5% by end-February 2003, some 0.4 percentage points below the equivalent unadjusted rate, but in line with the average survey expectation. Further out, the adjusted curve indicates that the expected peak in interbank rates was around 5.1%, some way below the market forward curve.

(1) Recursive estimates are generated by estimating the model from May 1993 through May 1995 and then by sequentially estimating the model with one more observation until the full sample period is used.
(2) A cubic spline is a mathematical technique for fitting a curve through points where the slope of the curve and the change in the slope are smooth everywhere.

Chart 11
Market and risk premia adjusted forward curves on 18 February 2002(a)

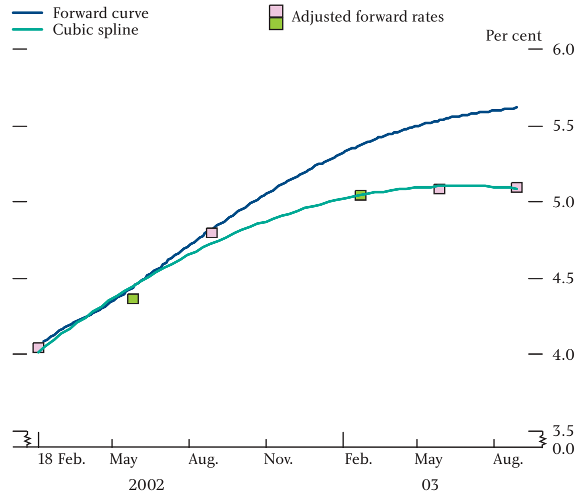

(a) The adjusted rates from the monthly model are depicted by green squares, and those from the quarterly model by pink squares.

Chart 12
Market and risk premia adjusted forward curves on 28 May 2004(a)

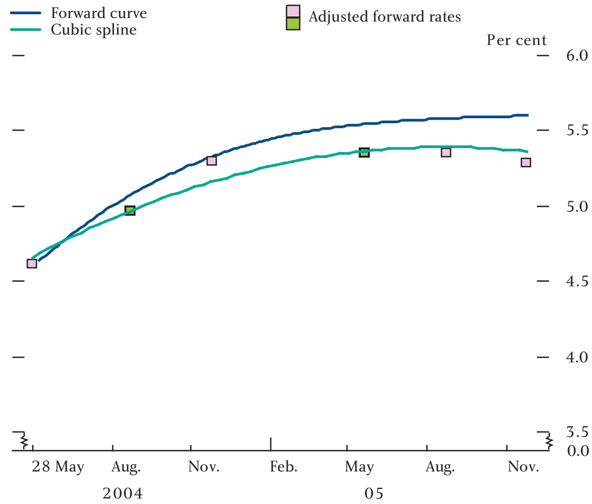

(a) The adjusted rates from the monthly model are depicted by green squares, and those from the quarterly model by pink squares.

So what of the more recent picture? Chart 12 shows that the interpretation is similar. The market forward curve has short-term interbank rates rising from a little over 4.6% on 28 May 2004 to around 5.6% by the end of next year. By contrast, the adjusted path shows a more gradual rise in interest rate expectations to a little under 5.3%.

Conclusion

Forward rates are perhaps the most common measure of expected future interest rates. But the presence of a risk premium can drive a wedge between forward rates and what the market expects future rates to be.

This article has used surveys of interest rate expectations to derive a measure of the risk premium. We find that the estimated risk premium implies a significant and time-varying difference between forward rates and survey expectations.

The survey data provide a *monthly* estimate of the risk premium and, therefore, expectations of future interest rates. But surveys are rarely timely, and with fast-moving financial markets they can become quickly out of date. Consequently, in this article we have also proposed a simple model of the survey-based risk premium using data available to us on a *daily* basis. This allows us to adjust the interbank forward curve on any particular day. To the extent that data observed on that day match historical experience, these adjusted curves should provide a closer reading of the market's expectation of the path of future short interbank interest rates than the forward curve alone.

Appendix

Testing the expectations hypothesis

To test the expectations hypothesis using survey data we use the following regression:

$$S_t(y_{t+n}) - y_t = a_0 + a_1{}^*(f_t^n - y_t) + e_t$$

where y_t denotes the three-month interbank rate at time t, $S_t(y_{t+n})$ denotes the survey expectation of this rate n periods ahead, f_t^n denotes the three-month interbank forward rate n periods ahead, and e_t is a zero-mean error term. If the expectations hypothesis holds then the constant a_0 and the slope coefficient a_1 will be insignificantly different from zero and one respectively (the null hypothesis).

Table 1 reports the estimates of a_0 and a_1 with standard errors in brackets.(1) The table also reports the regression correlation coefficient (R^2), and p-values for Wald coefficient restriction tests of the null hypothesis (χ^2). The sample period runs from May 1993 through April 2004. The table shows that, while the constants a_0 are insignificantly different from zero, at both horizons the slope coefficient a_1 is significantly different from its theoretical value under the null hypothesis. A joint test of the null hypothesis is decisively rejected with a p-value of less than 0.01. Thus, we reject the expectations hypothesis.

Table 1
Testing the expectations hypothesis

Horizon	a_0	a_1	R^2	χ^2
Three months	-0.00 (0.03)	0.42 (0.10)	0.49	0.00
Twelve months	0.04 (0.12)	0.51 (0.12)	0.65	0.00

Testing rational expectations

(a) Testing unbiasedness using survey data

To test unbiasedness we use the following regression:

$$y_{t+n} - y_t = b_0 + b_1{}^*[S_t(y_{t+n}) - y_t] + v_t$$

If survey expectations are rational then b_0 and b_1 will be insignificantly different from zero and one respectively.

Table 2 presents the estimation results over the full sample period. The table shows that at both horizons the constant b_0 and slope coefficient b_1 are insignificantly different from their theoretical values. A joint test of the null hypothesis also can not be rejected with a p-value of 0.47 at the three-month horizon and 0.13 at the one-year horizon. Consequently, we find that our survey expectations are unbiased.

Table 2
Testing rational expectations: unbiasedness

Horizon	b_0	b_1	R^2	χ^2
Three months	-0.07 (0.06)	0.97 (0.23)	0.21	0.47
One year	-0.40 (0.28)	0.79 (0.37)	0.15	0.13

(b) Testing efficiency using survey data

To test efficiency we use the following regression:

$$y_{t+\text{n}} - S_t(y_{t+n}) = c_0 + c_1{}^*y_t + c_2{}^*(f_t^n - y_t) + w_t$$

If survey expectations are efficient then all the parameters c_0, c_1 and c_2 will be insignificantly different from zero. Table 3 presents the estimation results over the full sample period. It shows that all parameters at both horizons are individually insignificantly different from zero. A joint test of the null hypothesis also cannot be rejected with a p-value of 0.64 at the three-month horizon, and 0.86 at the one-year horizon. Consequently, we find that our survey expectations are efficient.

Table 3
Testing rational expectations: efficiency

Horizon	c_0	c_1	c_2	R^2	χ^2
Three months	0.14 (0.64)	-0.05 (0.11)	0.43 (0.35)	0.17	0.64
One year	0.20 (9.93)	-0.12 (1.45)	0.03 (1.04)	0.02	0.86

Parameter estimates used to construct the adjusted forward curves

Table 4 reports the parameter estimates from the monthly model (shown in Table A) together with those from the quarterly model (in italics). Standard errors are

(1) Estimation is by OLS, with standard errors calculated using Hansen's (1982) generalised method of moments to correct for overlapping observations.

in brackets and significant variables are in bold. The sample period for the quarterly model is from May 1993 to March 2004. As the day of the survey changes from month to month, the horizon given in the table is to the nearest month.

Table 4
Modelling the survey-based risk premium

	3 months	*7 months*	12 months	*16 months*	*19 months*
Const.	-0.10 (0.08)	***-0.35*** ***(0.18)***	**-0.52** **(0.20)**	***-0.58*** ***(0.10)***	***-0.54*** ***(0.11)***
Slope	**0.26** **(0.03)**	***0.34*** ***(0.05)***	**0.42** **(0.06)**	***0.42*** ***(0.05)***	***0.46*** ***(0.04)***
Vol	0.03 (0.09)	*0.11* *(0.07)*	**0.44** **(0.21)**	***0.61*** ***(0.05)***	***0.60*** ***(0.06)***
CD	**1.01** **(0.52)**	***0.96*** ***(0.44)***	1.70 (1.17)	*-1.06* *(1.95)*	*-0.39* *(2.53)*
Swap	**0.16** **(0.08)**	*0.40* *(0.29)*	**0.57** **(0.14)**	***0.58*** ***(0.25)***	***0.65*** ***(0.30)***
R^2	0.74	*0.78*	0.72	*0.69*	*0.72*

References

Campbell, J and Cochrane, J (1999), 'By force of habit: a consumption-based explanation of aggregate stock market behaviour', *Journal of Political Economy*, Vol. 107, pages 205–51.

Campbell, J and Shiller, R (1991), 'Yield spreads and interest rate movements: a bird's eye view', *Review of Economic Studies*, Vol. 58, pages 495–514.

Consensus Economics, various issues.

Cortes, F (2003), 'Understanding and modelling swap spreads', *Bank of England Quarterly Bulletin*, Winter, pages 407–16.

Duffie, D and Singleton, K (1997), 'An econometric model of the term structure of interest-rate swap yields', *Journal of Finance*, Vol. 52, pages 1,287–321.

Fama, E and Bliss, R (1987), 'The information in long-maturity forward rates', *American Economic Review*, Vol. 77, pages 680–92.

Fama, E and French, K (1989), 'Business conditions and expected returns on stocks and bonds', *Journal of Financial Economics*, Vol. 25, pages 23–49.

Froot, K A (1989), 'New hope for the expectations hypothesis of the term structure of interest rates', *Journal of Finance*, Vol. 44, pages 283–305.

George, E (2002), Speech given to the Euromoney Bond Investors' Congress, London on 19 February 2002, and available at www.bankofengland.co.uk/speeches/speech162.pdf.

Hansen, L (1982), 'Large sample properties of generalised method of moments estimators', *Econometrica*, Vol. 50, pages 1,029–54.

Kamara, A (1994), 'Liquidity, taxes, and short-term Treasury yields', *Journal of Financial and Quantitative Analysis*, Vol. 29, pages 403–17.

Krishnamurthy, A (2002), 'The bond/old-bond spread', *Journal of Financial Economics*, Vol. 66, pages 463–506.

Liu, J, Longstaff, F and Mandell, R (2002), 'The market price of credit risk: an empirical analysis of interest rate swap spreads', *National Bureau of Economic Research Working Paper no. 8990*.

MacDonald, R and Macmillan, P (1994), 'On the expectations view of the term structure, term premia, and survey-based expectations', *Economic Journal*, Vol. 104, pages 1,070–86.

Maddala, G (1991), 'Survey data on expectations: what have we learnt?', in Nerlove, M (ed), *Issues in contemporary economics, macroeconomics, and econometrics: proceedings of the ninth world congress of the International Economic Association*, Macmillan, New York, pages 319–44.

Wachter, J (2004), 'A consumption-based model of the term structure of interest rates', *New York University Working Paper*, available at http://pages.stern.nyu.edu/~jwachter/research/hbinf2.pdf.