Fabio Fornari

+49 69 1344 8354
fabio.fornari@ecb.int

The rise and fall of US dollar interest rate volatility: evidence from swaptions[1]

Interest rate volatility, as implied by swaptions prices, rose in all major economic areas between 2001 and early 2004. The increase was particularly sharp for US rates and was more sizeable for short-term rates and swaptions with short expiration. Since the spring of 2004, US dollar volatilities have declined to the values recorded for euro rates and their term structure has flattened. The rise and fall of US dollar implied volatility reflected changes both in expectations of realised volatility and in the compensation for volatility risk.

JEL classification: G120, G130, G140.

The volatility of US dollar interest rates, as implied by the price of swaptions, increased substantially between 2001 and early 2004. The rise was much more marked than for euro rates, and increased especially for short-term rates and over short horizons, ie for swaptions with time-to-expiration of six months or less. However, both the higher average volatility of US rates and the relative peak in volatilities at short rates and horizons have receded considerably since the spring of 2004. As of end-March 2005, the volatility term structure was almost flat and the implied volatilities of US rates had fallen below those observed for the euro.

This special feature explores whether the rise in US dollar implied volatility was simply the counterpart of higher expected volatility, or whether it also reflected increased compensation for volatility risk. To investigate this issue we compare implied volatilities with forecasts of historical volatility derived from simulations of a GARCH model. We also seek to identify the main determinants of the gap between implied and forecast volatility, which is a metric of the compensation required for bearing volatility risk.

To anticipate the main results, a rise and fall in compensation for volatility risk has contributed significantly to the moves in US dollar implied volatilities. Compensation for volatility risk has usually been higher in the United States than in the euro area, particularly for short-term swap rates. After peaking in

[1] The views expressed are those of the author and do not necessarily reflect those of the BIS (where the article was written) or the ECB. The author wishes to thank Dimitrios Karampatos for research assistance and Claudio Borio, Frank Packer and Jakob Gyntelberg for useful discussions.

early 2003, it had retreated to roughly similar levels for both the United States and the euro area at end-March 2005. Among the determinants of compensation for volatility risk, the level of the interest rate and its volatility have had a sizeable positive effect. Positively sloped yield curves and negatively sloped volatility term structures are associated with lower compensation for volatility risk. Finally, macroeconomic surprises can affect compensation for volatility risk as well, though the size of these effects is quite variable.

The recent behaviour of volatility in swaption markets

Swaption markets provide an excellent opportunity to study the behaviour of implied interest rate volatility.[2] Compared to options on government bonds, swaptions are available on a larger set of interest rates (all the swap rates between one and 10 years) and for a broader spectrum of times to expiration (from one month to 10 years). This allows the construction of a term structure of implied volatilities for any given swap rate. Also, unlike options on government bonds or eurodeposits, swaptions have a constant life to expiration, which simplifies the empirical analysis.[3]

The implied volatilities used in this study are extracted from at-the-money swaptions on the US dollar and euro swap rates. They refer to the one-, five- and 10-year swap rates and are taken from swaptions with time-to-expiration of six months and two and five years. For both dollar and euro rates the sample analysed runs from 23 July 1997 to 30 March 2005.

Implied volatilities rose significantly from 2001, particularly for US rates

Implied interest rate volatilities are quite variable over time, and frequently differ substantially across currencies. From 1997 to the end of 2000, volatility was approximately equal for dollar and euro swap rates, generally moving in the 10–25% range. However, in 2001, implied volatilities on US rates started to rise well beyond those of euro swaps, and particularly sharply for short-term swap rates and options with short time-to-expiration (Graph 1). Though relatively subdued compared to US dollar volatility, the implied volatility for euro swaps also rose more for short-term rates and for short-expiration swaptions. Thus, the slope of the term structure of implied volatilities, ie the difference between long- and short-dated volatilities, became increasingly negative for both dollar and euro swaps.

Some large moves in implied volatilities were common to both dollar and euro rates ...

Some of the largest movements in the level of implied volatilities appear to be largely common to US dollar and euro rates. For instance, a jump in volatility occurred across currencies at the time of the 1998 Russian debt crisis and the failure of Long-Term Capital Management (LTCM) in the United States,

[2] Swaptions are options on swap rates. Since a swap rate is made up of a sequence of predetermined and stochastic payments related to future Libor rates, swaptions amount to options on the portfolio of such future Libor rates. They are priced according to the Black (1976) formula, where the volatility of the future swap rate plays the same role as the equity volatility in Black and Scholes (1973).

[3] As an example, a new three-month swaption on the 10-year rate is priced every day. By contrast, a newly issued three-month option on a 10-year bond has a decreasing maturity as it moves towards its expiration date.

Implied volatility of US dollar and euro swap rates

Annualised, in per cent

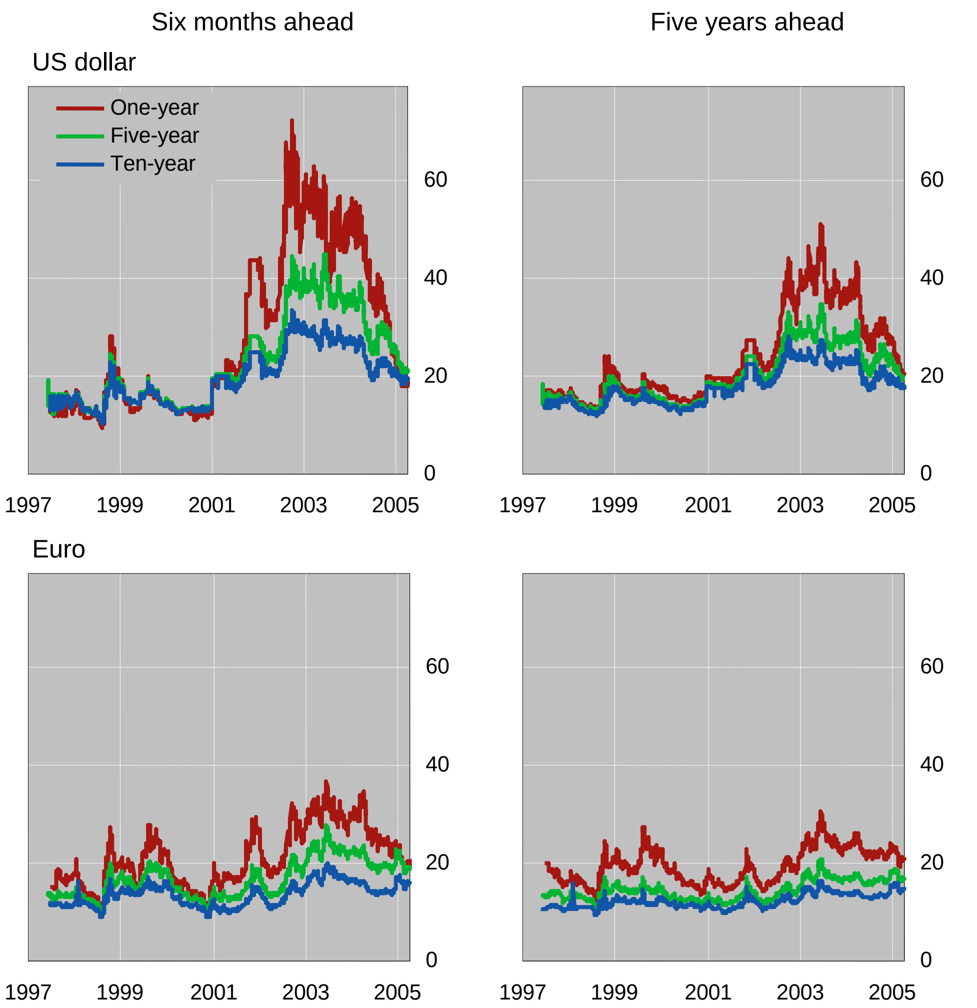

Source: Datastream.

Graph 1

when implied volatilities rose by approximately the same amount at different interest rate maturities and option expirations. The period of this rise was approximately the same length for both the dollar and euro. The jumps recorded at the end of 2000 and after the September 2001 terrorist attacks were also common to both dollar and euro swap rates, though the latter shock had a more pronounced impact on US dollar volatilities.

… while others were specific to the dollar

By contrast, the spikes in volatilities observed for US dollar yields from mid-2002 to early 2004 were generally unaccompanied by major movements in euro area volatilities. For example, between January 2002 and May 2004, the implied volatilities on the one-year US swap rate peaked at around 70% at the six-month horizon and 50% at the five-year horizon. The corresponding peaks for the implied volatilities of analogous euro swap rates were around 35% and 30%.

Monetary policy uncertainty may have contributed to dollar volatilities

It is likely that uncertainty about US monetary policy during the period, in particular the 2002–03 deflation scare, may have been partially responsible for an increase in anticipated volatility and hence may help to explain the gap between dollar and euro area implied volatilities.[4] Indeed, implied volatility was at high levels during the period, and was higher for short-term swap rates, which are more influenced by monetary policy. However, since the first tightening of Federal Reserve policy at the end of June 2004, markets appear to have been more certain about the path of interest rates. In this phase of diminishing monetary policy uncertainty, implied volatility has declined as target rate hikes have proceeded broadly in line with expectations of a gradual tightening (see BIS (2005, Chapter VI) for further discussion).

Did the increase in implied volatilities simply reflect expectations?

Implied volatilities should move with expectations …

Implied volatilities should clearly reflect economic agents' expectations about future volatilities over the interval spanned by the life of the option. However, since volatility changes through time in an unpredictable fashion, agents may also require compensation to bear volatility risk, ie the likelihood that future volatility deviates from its expected level. This compensation drives a wedge between implied and expected volatilities, which will be larger in a period of rapidly changing realised volatilities (ie when uncertainty about future volatility may be presumed to be highest).[5]

… though compensation for volatility risk may create a wedge

Did the rise in implied volatilities between 2001 and 2004 simply reflect a rise in the uncertainty about the future path of the swap rates or did it also reflect growing compensation required by market participants to bear that uncertainty? To answer this question, we must first specify a model for the behaviour of historical volatility which can generate volatility forecasts over various horizons. This will then be taken as a proxy for market participants' expectations of future volatilities. We assume that the historical interest rate volatility is well represented by an asymmetric GARCH model and estimate it for the logarithmic rates of change of the one- and five-year swap rates on expanding samples, all starting on 23 January 1997.[6] For each calendar day we use the features of the estimated model to simulate historical volatilities of

We model expected volatilities with an asymmetric GARCH model …

4 Differences in the level of interest rates in the two areas do not seem to account for the size of the gap recorded in implied volatilities. Admittedly, a rise in interest rate volatility – expressed as the standard deviation of the logarithmic rates of change of yields – could be expected in the context of falling interest rates, and interest rates fell more in the United States than in the euro area in 2002 and early 2003. But interest rates, whether at the short or long end, were not so far apart as to account for the volatility gap.

5 The compensation for volatility risk, which in the remainder of the article will be measured as the difference between the implied volatility and the expected realised volatility simulated through a GARCH model, is proportional to the relative risk aversion parameter of a power utility function. Bollerslev et al (2004), using data on the S&P 500 Index and its options between January 1990 and May 2004, find the coefficient of proportionality to be close to unity, so that the negative of the compensation for volatility risk equals the investor relative risk aversion (see also Bakshi and Kapadia (2003).

6 We use a GARCH scheme to estimate and forecast interest rate volatility since it is well known that this class of models is optimal even in the presence of various types of misspecification. On these issues see Nelson and Foster (1995) and, more recently, Fornari and Mele (2005). Jumps are not considered in this special feature. However, the typical average jump intensity and jump size identified for interest rates would not change the main message of our results. Our simulation methodology can be reconciled with Rosenberg and Engle (2002) and has been applied in a different fashion in Tarashev et al (2003).

Implied volatilities and confidence band for expected realised volatility of the US dollar and euro swap rates

Annualised; in per cent

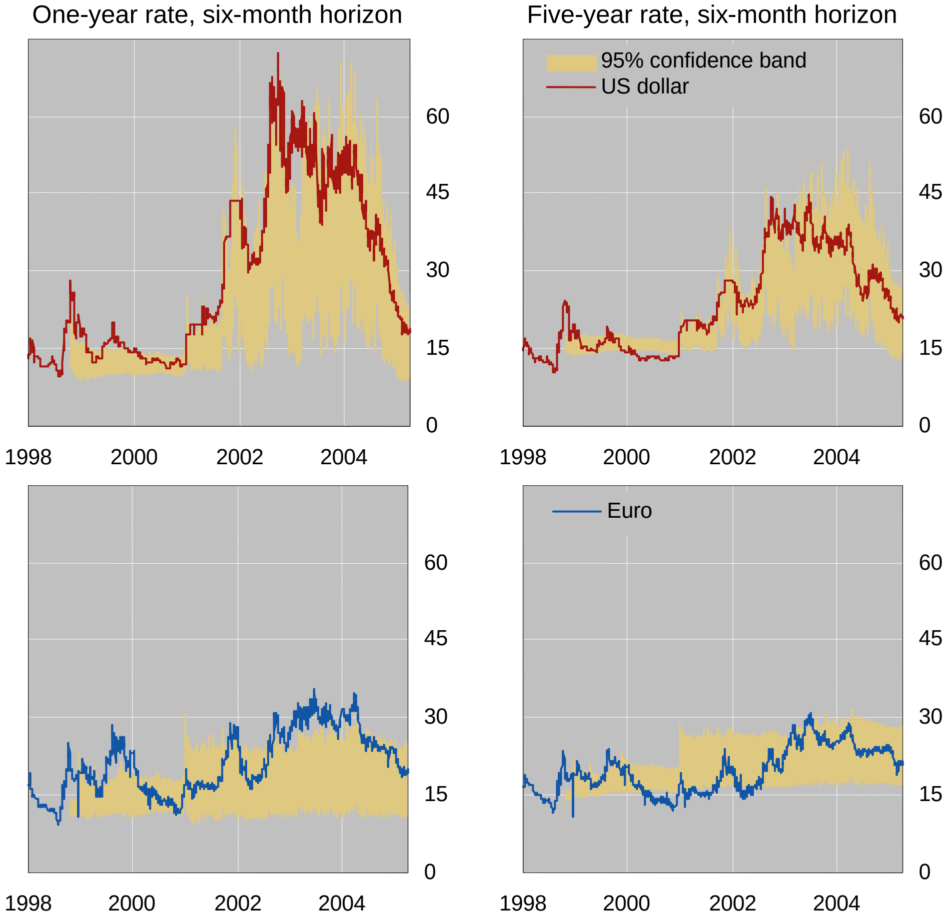

Sources: Bloomberg; BIS calculations.

Graph 2

the two interest rates (one- and five-year) over two forecast horizons (six and 24 months). For each of the two horizons and at each swap rate, forecast volatilities are then compared to implied volatilities. We define the compensation for volatility risk as the difference between the implied and average forecast volatilities.

... and use simulations to calculate confidence intervals

In principle, point estimates of the compensation for volatility risk would suffice for our analysis. However, we also use the simulations to calculate a probability distribution for the future expected volatility (see the box on page 93), whose percentiles provide a 95% confidence interval for the point estimates. The days in which the implied volatility lies outside the confidence interval can be assumed to represent periods of exceptionally high or low compensation for volatility risk. We look at the confidence interval so as to limit the risk of interpreting changes in our ability to estimate expected realised volatilities as changes in the compensation for volatility risk.

US implied volatilities moved above confidence interval bands in 2002 and 2003 ...

Especially in 2002 and 2003, implied volatilities were frequently and persistently outside the confidence interval of the forecasts for historical volatility (Graph 2). Episodes outside the confidence interval were rather common for one- and five-year dollar rates at both the six- and 24-month

Compensation for volatility risk of the US dollar and euro swap rates[1]

Annualised; in per cent

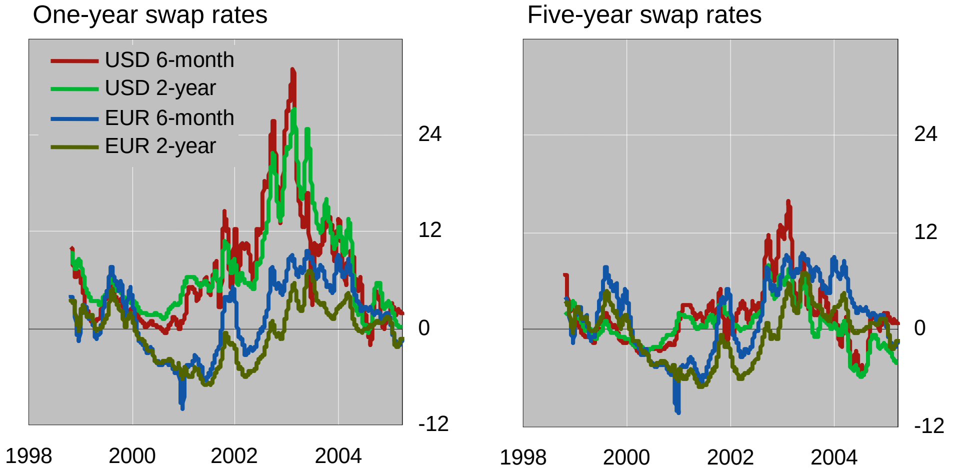

[1] The compensation for volatility risk is the difference between the implied volatility and the expected realised volatility; over the six-month and two-year horizons; one-month moving average.

Sources: Bloomberg; Datastream; BIS calculations. Graph 3

forecast horizons (Graph 2, upper panels, reports evidence for the six-month horizon). Results differ for euro rates, where implied volatilities move for an extended period above the upper limit of the confidence band only for short-term swap rates and, in fact, fall below the lower limit at times for longer-term rates (Graph 2, lower panels).

… consistent with high compensation for volatility risk especially at short horizons

Compensation for the volatility of dollar rates, as calculated by our model, has moved significantly through time and was exceptionally high for one-year swap rates at forecast horizons of six months, and between late 2001 and early 2004 (Graph 3). By contrast, for the one-year euro swap rate, the rise in volatility compensation in 2003 was much less sizeable than that for the analogous US rate. And for the five-year euro rate, in particular, compensation for volatility risk remained quite low, with implied volatilities less than forecast volatility being the rule rather than the exception.

While there was a noticeable correlation in volatility premia across countries and forecast horizons, co-movement was higher *within* countries than *across* countries. For both US swap rates (one-year and five-year), across the forecast horizons, the correlation was on average 0.9; for euro area rates, it averaged 0.6. By contrast, the US-euro area correlation was much lower, at 0.3 on average for both swap rates across the forecast horizons.

What determines compensation for volatility risk?

Compensation for volatility risk should be driven by the level of interest rates and implied volatility …

The obvious next question is what might explain the time variation in the compensation for volatility risk. According to standard finance theory, it should be related to the variables which influence the payoff of the derivative instrument. In the present application, we would thus expect the main determinants of the volatility risk premia to be the short-term interest rate level and its volatility.

Modelling historical volatility and generating volatility forecasts

We assume that historical interest rate volatility can be well represented by the following asymmetric GARCH(1,1) model (see Engle and Ng (1993)):

$$r_t = \mu + \phi \bullet r_{t-1} + \varepsilon_t$$
$$\varepsilon_t \mid I_{t-1} \sim N(0,\sigma_t^2)$$
$$\sigma_t^2 = \omega + \alpha \bullet \varepsilon_{t-1}^2 + \beta \bullet \sigma_{t-1}^2 + \gamma \bullet \max(0,-\varepsilon_{t-1})^2$$

where r_t denotes the logarithmic daily rates of change of a swap rate and σ_t^2 is its daily conditional variance; I_{t-1} is the information set, ie the past history of the interest rate series.

In a first step, the model was estimated for the one- and five-year swap rates of the United States and the euro area. To reproduce as closely as possible the expectations of economic agents at time t, the estimation was performed on expanding samples, the shortest of which starts on 23 January 1997 and ends on 15 October 1998 (450 daily observations). In this way volatility forecasts rely only on information available when forecasts were made. For each day we retain the parameters of the GARCH model, $\theta_t = (\mu,\phi,\omega,\alpha,\beta,\gamma)$, the time series of forecast errors (ε_t) and the historical volatilities (σ_t).

In a second step we use the information retained to produce, for each calendar day after 15 October 1998, forecasts of the historical volatility over various horizons. Each day we generate 2,000 future paths of the interest rate and its volatility, for each of the two interest rates (one- and five-year) and for two forecast horizons, six and 24 months. For each of these horizons we compute the expected volatility by averaging first across time-to-expiration[①] and finally across the 2,000 replications. This value is then compared, for each calendar day, to the implied volatility for the same swap rate and the same horizon. It is important to average volatility across time-to-expiration because implied volatility is an average volatility expected by a risk neutral investor over the life of the option.[②]

The structure of the simulation scheme is pretty much the same as the asymmetric GARCH(1,1) described above. The only difference is due to the distributional assumption placed on the standardised forecast errors ($\varepsilon_t/\sigma_t = z_t$). The implicit GARCH assumption that z_t are independently and identically normally distributed is rejected, due to the presence of asymmetry in excess of zero and kurtosis in excess of three. To reproduce these features we directly employ the estimated z_t in the simulation. For each calendar day, we randomly select an element of z_t and then loop over the following two equations, up to a two-year horizon:

$$\sigma_{t+1}^2 = \omega + \alpha \bullet (\sigma_t \bullet z_t)^2 + \beta \bullet \sigma_t^2 + \gamma \bullet \max(0,-(\sigma_t \bullet z_t))^2$$
$$r_{t+1} = \mu + \phi \bullet r_t + \sigma_{t+1} \bullet z_t$$

Given that in each calendar day we have 2,000 values for the expected volatility of each interest rate over the two forecast horizons, we can recover the distribution function of such expected volatilities. From this we calculate two measures of dispersion of the volatility forecasts, the standard deviation and the 2.5 and the 97.5 percentiles, both allowing us to build a confidence interval for the expected volatility.[③]

---

① As an example, in each working day, the expected six-month historical volatility is the average across replications (2,000) of the average volatility simulated in the six months after that specific day. ② Hence, comparing the time-t implied volatility to the time-t historical volatility (and not to the average historical volatility between t+1 and $t+\tau$, τ being the time-to-expiration of the swaption) defines a compensation for risk which rests on the assumption of a random walk in volatility, which has been strongly rejected by many applications of the GARCH methodology. ③ The volatility of volatility, ie uncertainty about future volatility, is the variable that should lead economic agents to adjust the swaptions price relative to what they would do by looking at the volatility forecast only. Swaptions are priced according to the Black (1976) model, whereby traders insert the expected volatility into a pricing formula similar to that of Black and Scholes (1973). However, given that volatility is time-varying, they will be more uncertain about this volatility when the volatility of volatility is very high and they will therefore increase the price of the swaption relative to the central forecast of the volatility when the confidence interval is wider.

We also include in our analysis the slope of both the term structure and the volatility term structure, as well as macroeconomic surprises. The term structure slope is a well known indicator of business cycle developments, and compensation for volatility risk may be influenced by the perceived stage of the business cycle. The slope of the volatility term structure, on the other hand, conveys information about the horizon over which interest rate uncertainty is concentrated. As for macroeconomic surprises, implied and forecast volatilities are likely to change significantly around the release of economic data, since economic agents map the size and sign of the surprise into the shape of the future distribution of interest rates, approximating the process according to which monetary authorities will react to such a surprise.[7]

… as well as macroeconomic surprises

Implied volatilities tend to fall after surprises, independently of their sign, since uncertainty is resolved after economic announcements are made (see Ederington and Lee (1993, 1996)). On the other hand, realised volatility, which is a major input into forecasting models, will always tend to rise after surprises. This suggests that our measures of risk compensation might be expected to fall upon a macroeconomic surprise. In general, we also divide surprises into positive and negative surprises, since the impact of surprises on interest rates and volatilities has been found to differ according to the sign of the surprise.

To shed light on the determinants of the compensation for volatility risk, we regress the premium, measured as the difference between the implied variance and the forecast historical variance (for the maturities and times-to-expiration reported in Graph 3), on the level of the three-month rate and a measure of interest rate implied volatility, on the slope of the yield curve, on the slope of the volatility term structure and on a set of economic surprises.[8]

Interest rates and volatility levels are indeed significant determinants ...

The results, as reported in Table 1, suggest that the premium required for the volatility risk of US rates has been mainly dependent on the level of the three-month rate and the level of short-term implied volatility. The slope of the term structure also appears to be a strong determinant of the compensation for risk, though the results are more variable. Its coefficient is negative in two out of the three significant cases, meaning that positively sloped term structures of interest rates tend to lead to lower compensation for volatility risk. Since

[7] Estimates of the effect of surprises on both interest rates and implied interest rate volatilities are reported in Fornari (2004). For a further discussion of the impact of macroeconomic surprises, see Fleming and Remolona (1999). Other variables originally selected, such as the credit spread (ie the yield differential between low-rated and high-rated bonds) and the swap spread (the differential between the swap rate and the government bond rate), were not statistically significant.

[8] The implied volatility of the one-year rate over a three-month horizon was employed in all regressions. All surprises are defined as the difference between the actual release of a variable and the consensus forecast and are then standardised to allow comparisons across types of news. The overall number of indicators is 35 (16 belonging to the United States, five to the euro area, seven to Italy, two to Germany and five to France). For the US economy: CPI, jobless claims, non-farm payrolls, durable goods orders, GDP, housing starts, Chicago PMI, index of leading indicators, PPI, retail sales, factory orders, capacity utilisation, industrial production, balance of trade, productivity. For the euro area: CPI, consumer confidence, industrial confidence, industrial production, PPI. For Germany: retail sales, IFO. For France: consumer confidence, consumer spending, CPI, industrial production, PPI. For Italy: advance CPI, industrial production, PPI, hourly wages, retail sales, business confidence, consumer confidence.

Regression of the compensation for volatility risk of US dollar swap yields on determinants				
	One-year swap rate		Five-year swap rate	
	Six-month horizon	Two-year horizon	Six-month horizon	Two-year horizon
Three-month rate	0.108 (9.1)	0.092 (9.5)	0.041 (5.9)	0.078 (17.9)
Implied	0.025 (11.8)	0.047 (26.0)	0.025 (16.5)	0.025 (21.5)
Slope	0.036 (2.4)	–0.042 (–3.5)	–0.032 (–3.5)	
Slope vol	0.016 (5.8)	0.041 (17.4)	0.025 (12.6)	0.020 (13.0)
ISM (+)	0.068 (2.6)	0.067 (2.5)	0.046 (2.1)	0.041 (2.7)
Jobless claims (+)			0.033 (2.6)	0.014 (1.8)
Industrial capacity (+)				0.064 (1.6)
Housing starts (–)	–0.132 (–3.4)	–0.082 (–2.0)	–0.064 (–2.2)	
Monetary policy		0.092 (2.0)		0.048 (2.6)

Note: The swap rate denotes the dependent variable of the regression, the horizon the time-to-expiration of the swaption. As an example, the combination of the one-year swap rate and the six-month horizon denotes the difference between the implied volatility of the one-year swap rate taken from a swaption whose time-to-expiration is six months and the corresponding forecast realised variance. The figures in parentheses are Student's t ratios. "Three-month rate" is the three-month eurodollar rate; "Implied" is the implied volatility of the one-year swap rate expected over a three-month horizon; "Slope" is the slope of the yield curve (10-year rate minus three-month rate); "Slope vol" is minus the slope of the term structure of volatility (volatility of the one-year rate minus the volatility of the 10-year rate, taken from swaptions with a three-month time-to-expiration); "ISM" is the index of supply managers; "Monetary policy" refers to dates on which Fed representatives gave speeches during the 2003 deflation scare. The symbol (+) or (–) after a macroeconomic surprise indicates that only the positive or the negative values of such a surprise have been used as regressors. The regression is run on daily data from 1 January 1999 to 8 April 2005. Table 1

positive slopes are indicative both of rising forward interest rates – which would command higher compensation for volatility risk – and of business cycle expansions – which might instead be expected to command lower compensation for volatility risk – the second component seems to have prevailed over the sample period. The slope of the volatility term structure has also had a negative impact on volatility risk compensation. When short-term expected volatilities are higher than long-term ones, which was typical of the sample analysed, risk compensation tends to rise.

Out of 32 surprises regarding US macroeconomic variables (16 variables split according to the sign of the surprise), three are found to influence risk compensation on days when the surprise is positive: the Index of Supply Managers (ISM), jobless claims and industrial capacity. Positive surprises tend to increase risk compensation. The only negative surprise which systematically affects the compensation for volatility risk is housing starts. A lower than expected figure for this variable is associated with a rise in the compensation.

By contrast, monetary policy events – FOMC meetings or speeches given by Federal Reserve representatives during the deflation scare period – have had only a mixed impact on the compensation for volatility risk, with the variable significant only intermittently, and the level of the regression coefficient implying a very limited economic effect.[9] These results suggest that the degree of monetary policy uncertainty may have driven implied volatilities and expectations of future volatilities in a similar fashion.

... though monetary policy events are not

When a similar regression is run for the euro area swap rate, we find that the coefficient of the three-month Euribor rate is negative, ie higher interest

Regression of the compensation for volatility risk of euro swap yields on determinants

	One-year swap rate		Five-year swap rate	
	Six-month horizon	Two-year horizon	Six-month horizon	Two-year horizon
Three-month rate	–1.06 (–2.8)	–2.98 (–10.2)	–2.32 (–14.6)	–1.76 (–6.1)
Implied	0.61 (5.5)	0.89 (11.8)	0.35 (9.0)	0.70 (10.3)
Slope	–3.95 (–7.1)	–3.79 (–7.9)	–2.57 (–14.2)	–5.83 (–15.6)
Slope vol	0.32 (2.2)	0.68 (6.9)	0.25 (4.8)	0.76 (7.9)
Italy retail sales (–)	–0.85 (–2.0)	1.08 (1.8)		0.90 (1.9)
Germany Ifo (+)	2.22 (1.9)	1.08 (2.1)		0.98 (2.9)
Italy PPI (+)		–2.10 (–3.8)		–1.76 (–6.1)
Euro area CPI (–)		–3.35 (–2.9)		0.70 (10.3)
Euro area PPI (–)	–2.83 (–3.4)	–1.45 (–2.7)		–5.83 (–15.6)
Italy PPI (–)		–2.98 (–10.2)		0.76 (7.9)
Euro area conf ind (+)		0.89 (11.8)		0.90 (1.9)

Note: The swap rate denotes the dependent variable of the regression, the horizon the time-to-expiration of the swaption. As an example, the combination of the one-year swap rate and the six-month horizon denotes the difference between the implied volatility of the one-year swap rate taken from a swaption whose time-to-expiration is six months and the corresponding forecast realised variance. The figures in parentheses are Student's t ratios. "Three-month rate" is the three-month euro rate; "Implied" is the implied volatility of the one-year swap rate expected over a three-month horizon; "Slope" is the slope of the yield curve (10-year rate minus three-month rate); "Slope vol" is minus the slope of the term structure of volatility (volatility of the one-year rate minus the volatility of the 10-year rate, taken from swaptions with a three-month time-to-expiration); "CPI" is the consumer price index; "PPI" is the producer price index; "conf ind" is the confidence index. The symbol (+) or (–) after a macroeconomic surprise indicates that only the positive or the negative values of such a surprise have been used as regressors. The regression is run on daily data from 1 January 1999 to 8 April 2005. Table 2

[9] For a list of episodes and the associated dates, see Bernanke et al (2004).

rates tended to lead to lower compensation for volatility risk (Table 2). This finding appears to be mostly driven by behaviour subsequent to 2000, when the Euribor rate moved in a much narrower range compared to the eurodollar rate, in a context of rising compensation for volatility risk in the euro area. The remaining financial variables (implied volatility, slope of the yield curve and slope of the volatility term structure) have the same sign as that observed for US dollar-based regressions, with the slope of the yield curve in fact exhibiting a more uniform and pronounced pattern. Unlike what has been observed for daily changes in interest rates, the compensation for volatility risk of euro yields does not seem to be driven more by US-specific news than by European news (Ehrmann et al (2005)). The macroeconomic variables found to be significant include both country-specific and euro area-wide surprises.

Conclusions

Interest rate volatility, as implied by swaptions, rose in all major economic areas from 2001 to early 2004, but particularly sharply for US rates at short-maturities and for short-expiration swaptions. We have analysed whether the rise in implied volatility was in line with expected volatility or was instead reflective of a significant increase in the compensation demanded for volatility risk. Our results suggest that between late 2001 and early 2004, dollar volatilities embodied a sizeable compensation for risk, which subsequently diminished considerably.

Compensation for volatility risk is mainly related the level of interest rates and volatility. Other variables, such as the slope of the term structure – which leads business cycle developments – and the slope of the volatility term structure – indicative of the horizon over which volatility is most pronounced – also affect the compensation. Positive macroeconomic surprises tend to lead to a rise in risk compensation as well. In contrast to the rise in implied volatilities, however, the rise in risk compensation does not appear to have been strongly affected by episodes of monetary policy uncertainty.

The compensation for volatility risk, as calculated, is of course dependent on the model employed to compute such forecasts. The models used in this paper do not take into account the possible presence of jumps in the interest rate process, nor more complex distributional assumptions for the forecast errors. An agenda for future research might be to explore the robustness of our findings to more general models.

References

Bakshi, G S and N Kapadia (2003): "Delta-hedged gains and the negative market volatility risk premium", *Review of Financial Studies*, 16, 527–66.

Bank for International Settlements (2005): *75th Annual Report*.

Bernanke, B S, V R Reinhart and B P Sack (2004): "Monetary policy alternatives at the zero bound: an empirical assessment", *Finance and Economics Discussion Series*, Federal Reserve Board, 2004–48.

Black, F (1976): “The pricing of commodity contracts”, *Journal of Financial Economics*, 3, 167–79.

Black, F and M S Scholes (1973): “The pricing of options and corporate liabilities”, *Journal of Political Economy*, 31, 637–54.

Bollerslev, T, M Gibson and H Zhou (2004): “Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities”, *Finance and Economics Discussion Series*, Federal Reserve Board, 2004–56.

Ederington, L H and J H Lee (1993): “How markets process information: news releases and volatility”, *Journal of Finance*, 48, 1161–91.

——— (1996): “The creation and resolution of market uncertainty: the impact of information releases on implied volatility”, *Journal of Financial and Quantitative Analysis*, 31, 513–39.

Ehrmann, M, M Fratzscher and R Rigobon (2005): “Stocks, bonds, money markets and exchange rates: measuring international financial transmission”, *NBER Working Paper Series*, No 11166, March.

Engle, R F and V K Ng (1993): “Measuring and testing the impact of news on volatility”, *Journal of Finance*, 48, 1749–78.

Fleming, M J and E M Remolona (1999): “What moves bond prices?”, *The Journal of Portfolio Management*, Summer, 28–38.

Fornari, F (2004): “Macroeconomic announcements and implied volatilities in swaption markets”, *BIS Quarterly Review*, September.

Fornari, F and A Mele (2005): “Approximating volatility diffusions with CEV-ARCH models”, *Journal of Economic Dynamics and Control*, forthcoming.

Nelson, D and D Foster (1995): “Filtering and forecasting with misspecified ARCH models II: making the right forecast with the wrong model”, *Journal of Econometrics*, 67, 303–35.

Rosenberg, J and R F Engle (2002): “Empirical pricing kernels”, *Journal of Financial Economics*, 64, 341–72.

Tarashev, N, K Tsatsaronis and D Karampatos (2003): “Investors’ attitude towards risk: what can we learn from options?”, *BIS Quarterly Review*, June.