Y.K. Chan
(212) 816-8337
yuenkwok.chan@ssmb.com

Hedging the Volatility Skew

Among the risk factors that determine mortgage prices are the quoted volatilities of options on swaps. A high level of implied volatility means a more expensive hedge for mortgages, hence a cheaper valuation. Large volatility exposure of a portfolio can be reduced with swaptions. However, mortgage portfolio managers sometimes intend to take a position in market volatilities, when they are historically high and are expected to decrease. At other times, the volatility exposure might be significant, but not high enough to justify the transaction costs and management efforts involved in the use of options. In these cases, it is useful to examine whether duration hedges with interest-rate benchmarks can be improved by exploiting the correlation of the volatility and rate levels, if such a correlation exists and can be predicted.

To illustrate, consider a mortgage-backed security with an effective duration of 5.0 and hedged with the ten-year Treasury note. For simplicity, assume that the Treasury note has a duration of 7.0. The hedge ratio is accordingly 5/7=0.71. Suppose the security has a volatility duration of 0.2. In other words, for every increase of 1% volatility the bond loses 0.2% of its value. Now, if we believe that every 100bp rally of the ten-year rate would be accompanied by an increase of 2% volatility, then it behooves us to adjust the hedging ratio to compensate for the anticipated price change resulting from the volatility change. Accordingly, the proper hedging ratio would be (5-2*0.2)/7=0.66.

This immediately raises a question: Is there any evidence to support the belief that rate moves have predictive power on volatility moves?

The answer seems to depend on the time period. In some years there was a very strong correlation — in other years, none. Figure 25 plots the quoted Black volatility of five-into-ten swaptions against the ten-year Treasury rate for the period since January 2000. Figures 26 and 27 are similar plots for the years 1999 and 1998, respectively.[9] The drastically different shapes of these plots clearly suggest that any estimation of correlation should be done in such a way that a change in regime is detected quickly.

Figure 25. Volatility of 5x10 Swaption Versus the Ten-Year Treasury Rate (Jan 00–Present)

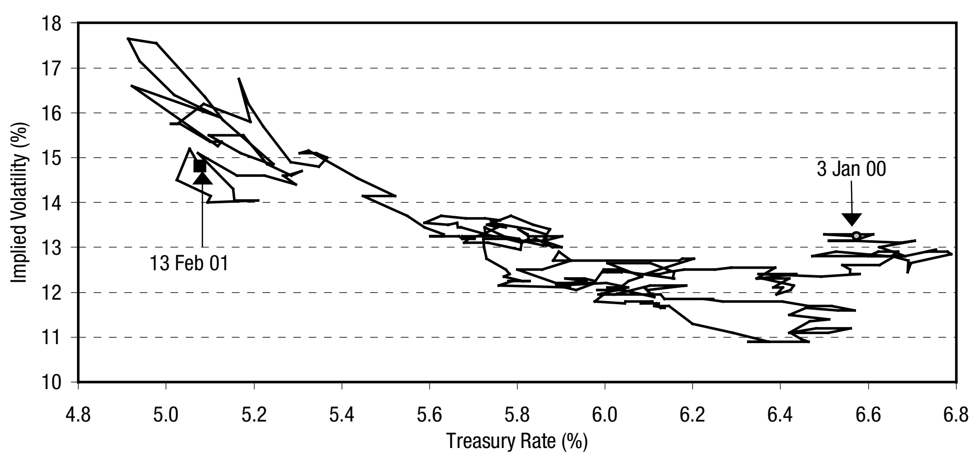

Source: Salomon Smith Barney.

[9] Despite appearances, these graphs have nothing to do with the Human Genome Project.

Figure 26. Volatility of 5x10 Swaption Versus the Ten-Year Treasury Rate (1999)

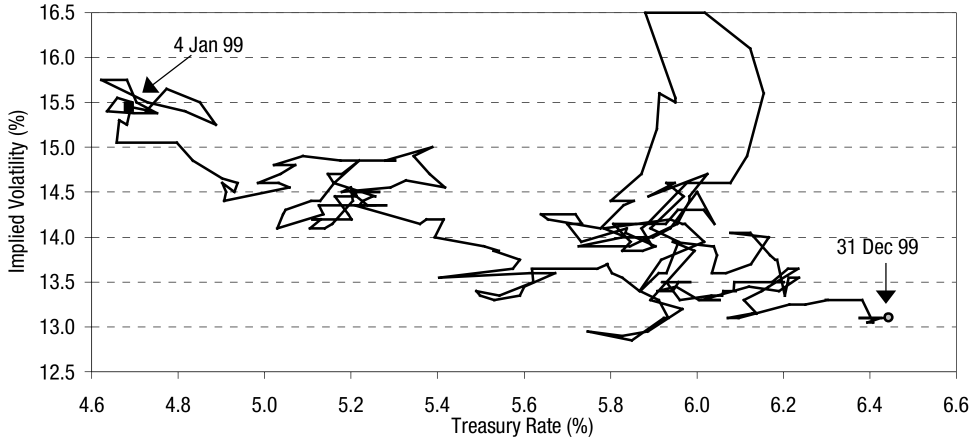

Source: Salomon Smith Barney.

Figure 27. Volatility of 5x10 Swaption Versus the Ten-Year Treasury Rate (1998)

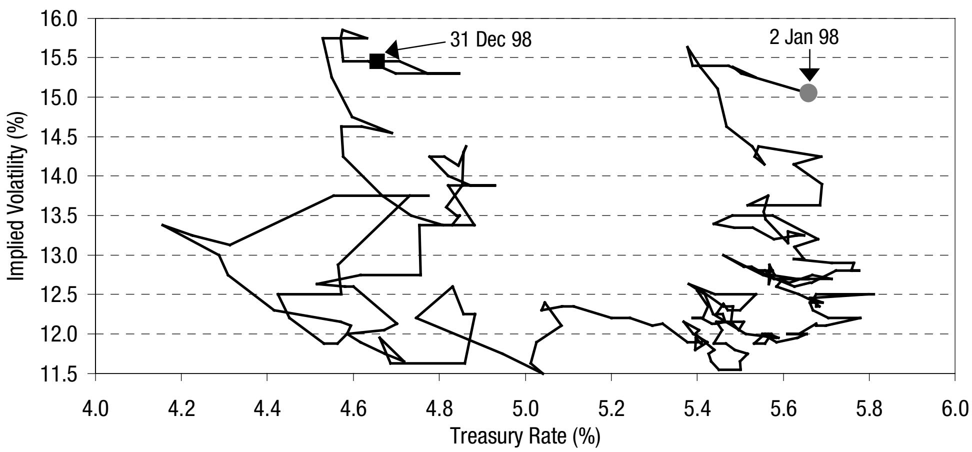

Source: Salomon Smith Barney.

Although there are sophisticated tools designed to predict volatility in a time series, here we test the performance of a simple rolling regression scheme. For each day since January 2, 1998, we regress the past moves in volatility on the corresponding moves in rate. We use exponential weighting with a mean decay time of three months, so that data in the distant past is given little weight. The choice of three months as the mean decay time is with the hope that it is long enough for a significant estimate of correlation, yet short enough to adjust to the nonstationary characteristics of the volatility series so prominent in Figures 25 through 27. The regression coefficient, commonly called the volatility skew in the current context, is used to predict the next day's volatility move, conditioned on the rate move. The standard deviation of the residual errors is compared to that of the actual moves. Both are tabulated in Figure 28.

Figure 28. Predicting Volatility Skew

Period	Average Skew; Vol Increase per 100bp Rally	Standard Deviation of Vol Moves	Standard Deviation of Residual Error
1998	0.90	0.18	0.18
1999	-0.03	0.16	0.16
2000–Present	0.52	0.23	0.19
Feb 13, 2001	3.90	-	-

Source: Salomon Smith Barney.

Mirroring Figures 25 through 27, the skew coefficient is far from stationary. Fortunately, the exponentially weighted scheme, with three-month decay time, at least seems to limit the error from an obsolete skew so that the residual error is no larger in standard deviation than the actual moves. In the past year, there is actually noticeable improvement in using the prediction. The estimated skew coefficient today, 3.9% volatility per 100bp rally, is significantly higher than indicated by quotes on out-of-the-money constant-maturity options on swaps, the latter translating to approximately 2.10.

What should the investor do? Quotes for out-of-the-money and in-the-money options may be a good indication of the skew. However, in that regard there is a question as to whether the market quotes contain a large part of risk premium, which does not reflect expectation. Market quotes are important information for derivatives pricing, but not necessarily for anticipating moves. As an alternative to quotes for out-of-the-money and in-the-money options, the rolling regression seems usable. Whether either proves useful for the next 100bp rally, only time can tell.

In the meantime, the above idea can be applied to other risk factors, as has been done by various researchers. The more immediate risk factors are the other key rates: the two-, five- and 30-year rates. Others factors include swap spreads, the current-coupon OAS, and, for investors in CMO derivatives, the OAS of IOs and POs. Given a move in the ten-year rate, a good prediction of these other risk factors could improve on the effective-duration hedge (in a manner similar to the adjustment for the volatility skew mentioned previously). Generally, given moves in the several key Treasury rates, a good prediction of the remaining risk factors would improve the performance of partial-duration hedges. Indeed, the adjustment for volatility skew should be done in conjunction with all the other risk factors. We will discuss this in a future article.