EC372 Economics of Bond and Derivatives Markets

Volatility

1. What is volatility and why does it matter?

In the context of this note, *volatility* means the 'standard deviation of the rate of return on an asset'. The asset will be an company's stock or a share price index, the application being to the *S&P 500* index. The *variance* is defined as volatility-squared, i.e., the 'squared standard deviation of the rate of return on an asset'. Volatility and variance have a time dimension: conventionally the unit is one year, and values are quoted as percentages.

Volatility and variance are thus indices of the variability of an asset's rate of return. Volatility matters because it is a determinant of many financial decisions and hence of security prices. For example: mean-variance portfolio selection requires the variances (and covariances) of asset returns; option price formulæ are functions of, among other variables, the underlying asset's volatility. More generally, but less precisely, volatility reflects the stability or otherwise of asset prices – volatility increases at times of unease in financial markets, dramatically so during crises.

2. Explicit versus implicit volatility

By construction volatility is a 'summary statistic' based on the probability distribution of an asset's rate of return. But what is the source of the probability distribution? In finance, probabilities appear in two distinct ways: (i) as a way of expressing investors' beliefs about asset prices, and hence rates of return – call these 'belief' probabilities[1]; (ii) as a reflection of the absence of arbitrage opportunities, i.e. 'risk-neutral' or 'martingale equivalent' probabilities. Remember that the absence of arbitrage opportunities is *equivalent* to the existence of risk-neutral probabilities.

In principle, measures of volatility can be 'backward-looking' (what volatility *was*) or 'forward-looking' (what volatility is forecast *to be*). Although types (i) and (ii) could be either, measures assumed to be based on belief probabilities are almost always backward-looking, those for risk-neutral volatility being constructed as forward-looking.

While there exist circumstances in which the belief probabilities would result in the same value of volatility (and variance) as the risk-neutral probabilities, there is no guarantee that the required circumstances hold. Indeed, empirical evidence suggests that there may be significant differences between the two.[2] Hence, a distinction should be made (but is often overlooked) between volatility obtained from belief probabilities and risk-neutral probabilities.

[1] Sometimes that the belief probabilities are referred to as 'objective' probabililites, but this is misleading because it suggests that there are 'true' probabilities that somehow exist separately from and independent of investors' beliefs.

[2] See Bollerslev, T. and H. Zhou, "Expected Stock Returns and Variance Risk Premia", *Finance and Economics Discussion Series, Federal Reserve Board, Washington, DC*, (unpublished, 2007-11), for a careful study of the 'variance risk premium', i.e., the difference between implied and realized variances.

In practice, the empirical counterparts of these two concepts of volatility are expressed as 'explicit' and 'implicit' volatility. Explicit volatility, used to reflect belief probabilities, is also known as 'realized', 'observed' or 'historical' volatility. The estimates of explicit volatility are based on statistical inferences from past rates of return (i.e., from realized asset prices).

Implicit volatility, used to reflect risk-neutral probabilities, is estimated from derivatives' prices in relation to their underlying asset prices: almost always the derivatives are European style options.

3. How is explicit volatility measured?

Explicit volatility is estimated as the standard deviation of past data on asset prices. The simplest method is to calculate the sample variance of proportional changes in an asset's price. Given that the interval of time between successive observations of the asset price is short (typically a day or less), the rate of return is dominated by the proportional rate of change in price, i.e., the capital gain or loss (the dividend per day, or less, is so small that it can be ignored).

Let S_t denote the asset's price at time t and let σ denote volatility. Usually, g_{t+1} is defined as $g_{t+1} \equiv (S_{t+1} - S_t)/S_t$, though in practice the continuously compounded rate of change of S_t is used instead, i.e., $g_{t+1} \equiv \ln[S_{t+1}/S_t] = \ln(S_{t+1}) - \ln(S_t)$. The standard deviation of g_{t+1} can now be used to estimate σ.

Suppose that a sample of daily asset prices is available, $S_0, S_1, S_2, \ldots, S_N$. The rate of return for each date is then given by $g_1, g_2, g_3, \ldots, g_N$, where $g_1 = \ln[S_1/S_0]$, $g_2 = \ln[S_2/S_1]$, $g_3 = \ln[S_3/S_2]$, and so on. Now σ^2 can be estimated by the sample variance:

$$\widehat{\sigma}^2 = \frac{(g_1 - \overline{g})^2 + (g_2 - \overline{g})^2 + (g_3 - \overline{g})^2 + \cdots + (g_N - \overline{g})^2}{N-1}$$

where $\overline{g}$ denotes the sample average, $\overline{g} = \sum_1^N g_t/N$. (Given the short time interval between t and $t+1$, $\overline{g}$ is likely to be tiny and is commonly ignored.) The estimate of volatility, $\widehat{\sigma}$, is just the positive square-root of the variance.[3]

Although the circumflex, $\widehat{\ }$, over σ is omitted from now on, it should be remembered that the value obtained depends on a sample of data – its 'true' value can never be inferred from a sample of data.[4]

So far it has been assumed (albeit not explicitly) that volatility is *constant* both across time and with respect to other potential influences on its value (e.g., 'investor sentiment'). Although this is a restrictive assumption, there is no consensus about how non-constant volatility should be modelled. In practice the methods used (some of them highly sophisticated) try to infer patterns of volatility that fit past data most closely, with little (if any) theoretical justification for the empirical specification chosen.

[3] As the time dimension for σ is normally chosen as a year, it will be necessary to multiply by a scaling factor that depends on the observation interval. See *Economics of Financial Markets*, page 484, for the details.

[4] 'True' in this context refers to the value of σ for the probability distribution from which the data on asset prices is generated. As this probability distribution is always purely hypothetical, the most that could be claimed is that $\widehat{\sigma}$ has desirable properties – essentially that it is close to σ in a precisely defined sense – *if* the sample of data were generated from the assumed probability distribution.

4. How is implicit volatility measured?

Begin by assuming a relationship (formula) between the price of a derivative (a European call or put option) with a set of variables, including volatility:

$$c = f(S, X, \tau, R, \sigma) \tag{1}$$

where c is the premium for a European call option, S is the underlying asset price, X is the exercise (strike) price, τ is the time to expiry, and R is the interest factor (typically $R \equiv e^{r\tau}$, r denoting a constant, risk-free interest rate).

If a specific functional form is assumed for the function $f()$ in (1), then with observations on c, X, τ, R it is possible to calculate (or 'back out') a value for σ. While it is typically not possible to obtain an explicit solution for σ as a function of c, S, X, τ, R, iterative methods are available to obtain numerical solutions.

Note that c, S, τ and R vary over time during the life of the option. By definition, the option contract specifies X. Volatility, σ, is also commonly assumed fixed, though whether it is appropriate to do so is highly problematical.

Early estimates of implicit volatility were obtained by applying the Black-Scholes model to provide the functional form for $f()$ in (1). This approach has its attractions, including (a) familiarity of the Black-Scholes model, (b) the assumptions of the Black-Scholes model imply that explicit and implicit variances (squared volatility) are equal.

But there are also damaging weaknesses, most importantly that different exercise prices (X values) result in different estimates of σ (for the same values of S, τ and R). Being a property of the probability distribution of the underlying asset's return, the value of σ should not vary with X, a property commonly stated with the assertion that there should be a 'flat smile'. However, under the assumptions of the Black-Scholes model, empirical evidence suggests that the 'smile' is *not* flat.

Also, even with a 'flat smile', if volatility varies across time, values calculated from the Black-Scholes equation will be biased.[5]

In view of the evidence against the Black-Scholes model (in particular, non-flat smiles), estimates of implicit volatility relying on weaker assumptions – so called 'model-free' measures – have been devised. The most well known of these is the Chicago Board Options Exchange (CBOE) 'VIX' index.

5. The VIX index

The CBOE VIX index provides an estimate of implicit volatility over a 30-day period from the present for the S&P 500 index of US stocks. The volatility measure, based on the prices of European call and put options on the S&P 500 index, relies only on the absence of arbitrage opportunities in frictionless markets and the assumption that the underlying asset price distribution is continuous

[5] For a rigorous analysis, see Britten-Jones, M. and A. Neuberger "Option Prices, Implied Price Processes and Stochastic Volatility", *Journal of Finance*, vol 55 (2), April 2000, pp. 839–866.

(i.e., does not exhibit discrete 'jumps'). The VIX index equals $100 \times \sigma$ for σ^2 calculated according to:

$$\sigma^2 = \frac{2}{T} \sum_i \frac{\Delta X_i}{X_i^2} \cdot R \cdot Q(X_i, T) - \frac{1}{T} \left[\frac{F}{X_0} - 1 \right]^2$$

where

T is the time to expiry;

$R \equiv e^{rT}$ is the risk-free interest factor;

$F \equiv R \cdot S$ is the forward price of the S&P 500 index, S, with 'delivery' at T, calculated from the put-call parity relationship: $F = R \cdot S = X + R \cdot (c - p)$, where c and p are call and put option premia for options with exercise price X (chosen as close as possible to the underlying asset price S);

X_i denotes exercise price for option i, a call if $X_i > F$, a put if $X_i < F$, and X_0 is the exercise price nearest below F;

ΔX_i measures the interval between exercise prices adjacent to i: $\Delta X_i = \dfrac{X_{i+1} - X_{t-1}}{2}$;

$Q(X_i, T)$ is the option premium corresponding to X_i with time T to expiry (more precisely, $Q(X_i, T)$ is the mid-point of the option's bid-ask spread).

Constructed in this way, each σ^2 can be interpreted as the market value of a portfolio of options with different exercise prices and the same expiry date. For each trading day, the value of σ^2 is obtained as a weighted average for two expiry dates – the nearest before and after 30 days from the present – the weights being proportional to the difference between the expiry dates and 30 days.[6]

Computed values of the VIX index since 1998 are shown in figure 1 on page 5. (data source: CBOE website).

6. Uses for the VIX index, and other measures of volatility

The most obvious use for implicit volatility is to provide a forward-looking indicator of anticipated variability in asset prices. Thus, the VIX index can be interpreted as a market forecast of variability in the S&P 500 index over the coming 30 days. It is a 'market forecast' in the sense that the VIX index is calculated from option prices on the S&P 500 index, hence reflecting the beliefs and risk-preferences of investors in making their decisions to hold or write options on the index. It should not be surprising, then, that the VIX index is referred to as the investors' 'fear gauge' – it reflects concerns about the variability of stock prices in the coming month.

The VIX index has several specific uses:

1. Futures contracts: the VIX index forms the underlying asset for futures contracts traded on the CBOE.

[6] A detailed example of the calculation is provided on the CBOE website at:

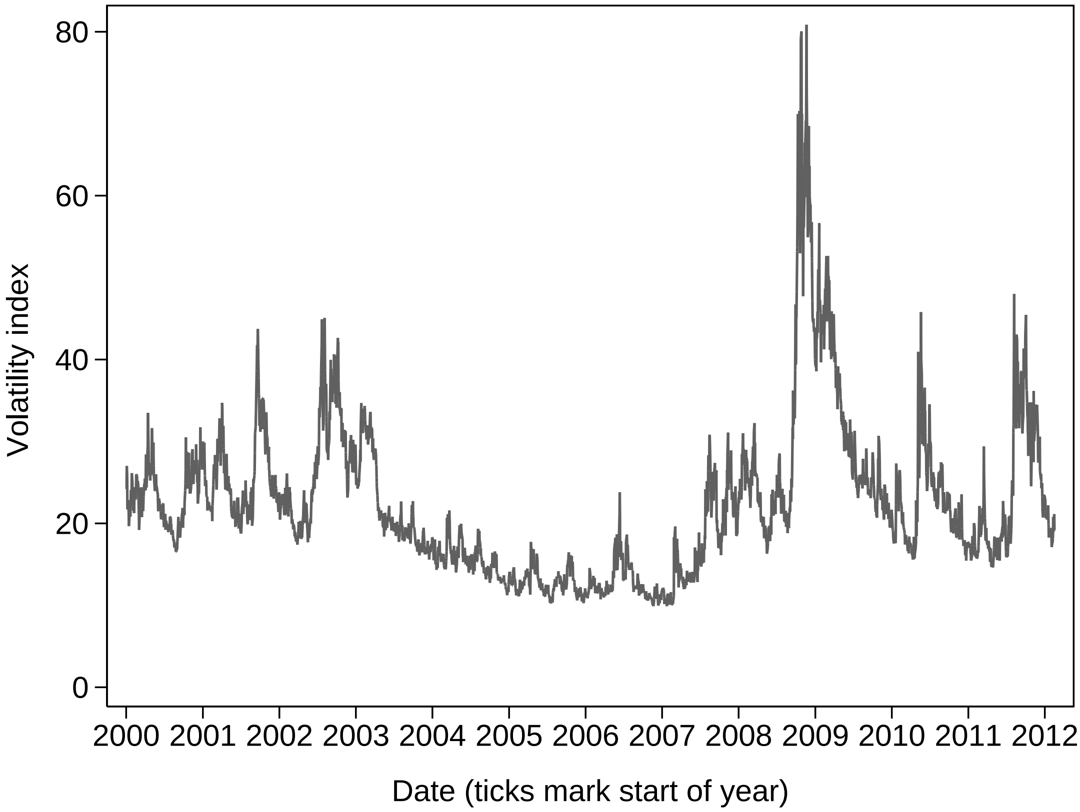

Figure 1: The VIX index: implicit volatility of the S&P 500 index

2. Options contracts: the VIX index forms the underlying asset for options contracts traded on the CBOE.

3. Swap valuation: 'variance swaps' are traded (over-the-counter) on the variance of the S&P 500 index over an interval from the present, t, to a specified date in the future, T. At T, the payoff on the swap is proportional to the difference between the realized variance of the S&P 500, from t to T, and a number, say X, agreed at the outset. The VIX index can be used to 'price' such swaps, i.e., to provide a guide for the value of X such that both parties to the swap accept the agreement. Variance swaps are thus a variety of forward contract.

4. Forecasting: the VIX index can be used as a market forecast of asset price variances during the ensuing 30 days.

5. Variance risk premia: the VIX index can be used to estimate variance risk premia (difference between implied and realized variance). Variance risk premia can then used in empirical studies of fluctuations in asset returns (see footnote 2 on page 1).

7. Recent fluctuations in the VIX index

Figure 2 shows the VIX index since 1st January 2008. Note the dramatic increase beginning in September 2008, peaking in October and November, then declining to fluctuate around a level high by comparison with months prior to the onset of the credit crisis in August/September 2008.

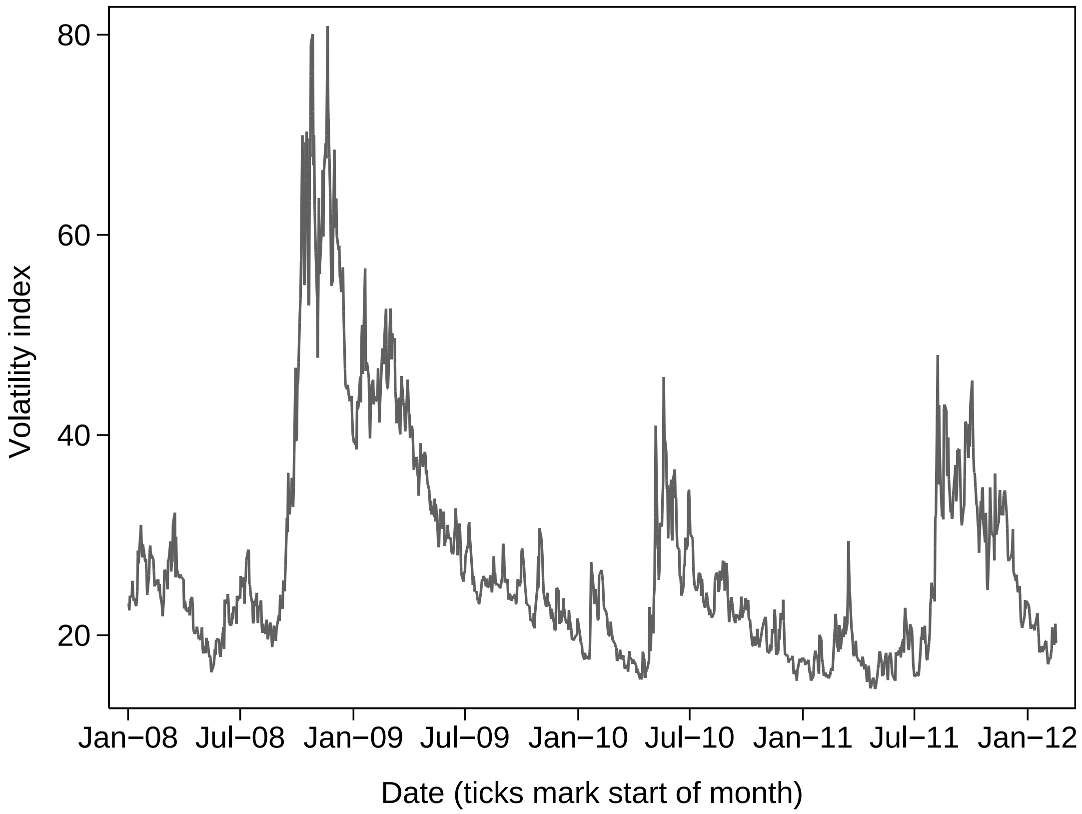

Figure 2: The VIX index in 2008–12

*****