



Don't be fooled by the surprises

Nicolas Mousseot and Florent Brunes of BNP Paribas examine the impact of earnings announcements on volatility

It is well known that stock prices can react strongly to quarterly or annual earnings announcements, as this is an opportunity for the firm to provide more information regarding its past and expected future performances. However, little is said regarding realised and implied volatility trends around the reporting season. We, therefore, discuss how implied and realised volatilities evolve around earnings announcements. We further suggest that the best trading strategy may be to sell implied volatility after the earnings date, particularly after the company has disappointed the market, as implied volatility tends to revert back to its normal level too slowly.

2005: a year full of surprises?

The effect of earnings announcements on volatility is of particular interest this year for several reasons:

- Implied and realised stock volatilities are so low that a variation of 1% or 2% in the stock price may have a great impact on volatility.
- In Europe in particular, companies have been recently asked to switch to the new IAS/IFRS accounting standards, which may yield more surprises, as the effect of the introduction of the new accounting rules is difficult to estimate.

We are thus in a particular environment where volatility may vary more than usual around the reporting season.

Volatility around earnings dates

Let us assume that the Black-Scholes hypothesis holds. In such a case, implied volatility can be viewed as the average of future realised instantaneous volatility. Implied volatility should, therefore, be high – provided that the announcement date is before the option's maturity date and then drops on the day earnings are released.

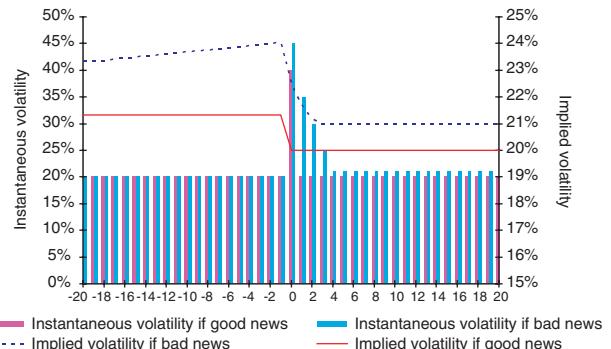
As Black highlighted, the dynamics of realised volatility may differ, depending on whether the news is good or bad. Realised volatility will take a few days to revert back to its normal level, and is higher for bad than for good news.

Since implied volatility is, broadly speaking, the average of the expected instantaneous volatilities, it should thus increase until the announcement date and then decrease in a few days (see graph 1).

By the way, what is a surprise?

A surprise is, in general, modelled by either looking at the stock price performance on the day of the announcement, or by comparing the differ-

Graph 1: Theoretical evolution of implied volatility around earnings announcements



Source: BNP Paribas

ence between the analysts' consensus for earnings some time before the announcement, say one month, and the actual announced earnings.

The stock price performance can be simply measured by subtracting the market return from the stock return, taking into account the stock's beta:

$$Surprise = R_i - \beta_i R_M \begin{cases} \text{positive if } >0 \\ \text{positive if } <0 \end{cases}$$

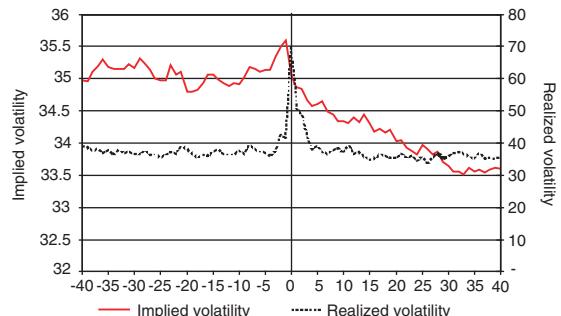
Another approach to defining surprise is to look at the difference between the earnings consensus some time before the announcement and the realised earnings. However, dealing with an earnings consensus may yield biased outcomes. First, analysts have historically been overly optimistic and, hence, their earnings estimates may be upwardly biased. Second, earnings estimates may be less accurate in the coming quarters or years due to changes in accounting standards. For example, in Europe, the effects of the new IAS/IFRS accounting rules imply changes in calculating projects, pensions, derivatives, etc, and their implications are not yet entirely clear.

We thus chose market-adjusted stock return as an objective measure of earnings surprises.

Overall effect of a surprise on implied volatility

Our universe consists of 249 stocks within the DJ Stoxx 600 index for which there are listed options, spanning over eight years of data and com-

Graph 2: Evolution of realised and implied volatility around earnings announcements



Source: BNP Paribas, Bloomberg

prising 3,795 earnings dates. We first assess whether historical and implied volatilities evolve as predicted by the theory. Graph 2 shows the average implied volatility trend of the 249 stocks included in our universe 40 days before and after their earnings dates. Six-week implied volatility tends to increase by 0.5 point of volatility in the days ahead of the earnings date and then decrease by more than two points of volatility between its peak – the day prior to the announcement – and trough – several weeks after the announcement.

Realised volatility also evolves as predicted. Its average level is more or less constant over time, with the exception of the day of the announcement, when it rises to a level twice as high as usual.

Surprises, surprises: good and bad

The next step is to analyse whether changes in implied volatility around earnings announcements are related to the sign of the surprise.

Graph 3 shows that the surprise effect is asymmetric: while a positive surprise is followed by a significant decline in implied volatility, implied volatility stays at a relatively high level if the company disappoints investors.

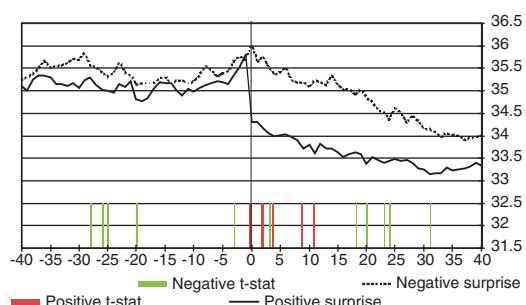
When is it best to buy/sell implied volatility?

The previous analysis shows that, in general, implied and realised volatility trades in accordance with what the theory predicts.

However, we have previously observed that, when the company disappoints the market, implied volatility tends to stick to a relatively high level after the earnings date.

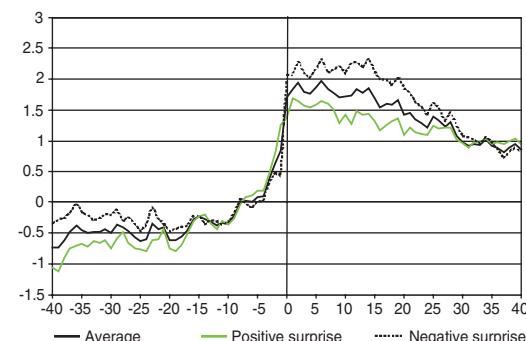
Graph 4 reports the difference between the implied volatility and stock

Graph 3: Effect of positive and negative surprises on implied volatility



Source: BNP Paribas, Bloomberg

Graph 4: Difference between implied and future realised volatilities



volatility realised during the life of the option. One can note that on average the difference between implied and realised volatility is equal to around two points just after earnings are released. The difference is even wider in the case of bad news, which means that implied volatility remains too high after bad news is released.

Therefore, one systematic volatility strategy to benefit from this anomaly would be to sell straddles after the announcement and delta-hedge the short straddle position in order to realise the difference between implied and realised volatility.

Conclusion

As hedge funds and proprietary trading desks are looking for catalysts for a potential increase in volatility, earnings announcements will probably be a focus of attention in 2005. More immediately, the introduction of the new European accounting standards in January is likely to further increase uncertainty around earnings announcements. The empirical analysis of the dynamics of implied and realised volatilities appears to be in line with what the theory predicts, as implied volatility tends to rise slowly before earnings announcements and decline sharply after, while realised volatility experiences a peak on the day of the announcement.

Without further information, we advise investors to wait until after the earnings announcement to take option positions. In particular, we suggest selling straddles when the company disappoints the market. ■

CONTACTS

Europe

Philippe Dardier

T. +33(0)1 40 14 28 97

E. philippe.dardier@bnpparibas.com

Europe

Talbot Stark

T. +44 (0)20 7595 8349

E. talbot.stark@uk.bnpparibas.com

USA

Nicholas Kello

T. +1 (212) 841 3796

E. nicholas.kello@americas.bnpparibas.com

Japan

Mark Arnell

T. +(813) 5290 8343

E. mark.arnell@japan.bnpparibas.com

Hong-Kong

Emanuel Breiter

T. +852 2978 3656

E. emanuel.breiter@asia.bnpparibas.com