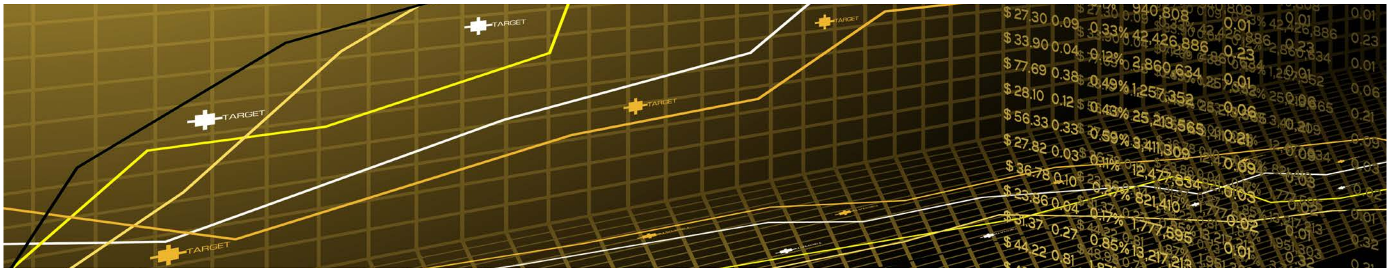

TIME AND VOLATILITY IN E-MINI S&P 500 FUTURES OPTIONS: PART II

BY LAWRENCE MORGAN
14 AUGUST 2013

I. INTRODUCTION

In the first article of this series on E-mini S&P 500 Options, I noted that "Options are a standard piece of hardware in the 'tool-box' of derivatives traders – whether they specialize in derivatives operations or manage funds more comprehensively as investors, speculators, or traders"; that the E-mini S&P 500 market is one of the major benchmarks of the U.S. equity market; and that it is rewarding to look in detail at the option market based on that futures contract.

That first article treated opportunities that come from the *time* aspect of options. This one will concentrate in the same way on *volatility.* This aspect of option valuation brings unique opportunities to options trading. As readers are aware, option volatility is one of the most studied aspects of derivatives markets – learned articles, traders' blogs, entire books have been written about it. This note can only outline some of the basic aspects of incorporating the concepts of volatility into options strategies.

II. HISTORICAL AND IMPLIED VOLATILITY

Option traders and analysts typically refer to the historical volatility (HV) (also called "realized volatility") and implied volatility (IV) of a market. Historical volatility is simply the standard deviation of returns based on the underlying price of a stock, commodity, index, bond, etc. That is, the calculated standard deviation of

$$\frac{P_t}{P_{t-1}} - 1$$

This is annualized so that it appears as X% per year. HV is usually based on daily price data, but it can be any period one wants to examine, and the standard deviation is frequently calculated for periods of various lengths, such as 10-day, 30-day, or 50-day volatility.

Option strategists are primarily interested in implied volatility, which is supposedly *the market's view of what realized volatility will be in the future*, that is, what volatility will be realized during the period until expiration of whatever option you're looking at. No one knows the "market's view" of volatility, of course, but it is "implied" in option prices. Volatility is one of the risk factors in an option pricing formula such as the celebrated Black-Scholes formula, appearing as the Greek letter sigma (σ) to indicate the standard deviation of the returns. Consequently, if you know the other risk factors (underlying price, exercise price, time to expiration, interest rate) and the option premium you can solve for the volatility parameter – therefore it is "implied" in premium. When someone says volatility is rising or falling, or that one wants to "buy" or "sell" volatility, or when one analyzes how an option premium will change in reaction to a change in volatility, this is the volatility the trader or strategist has in mind.

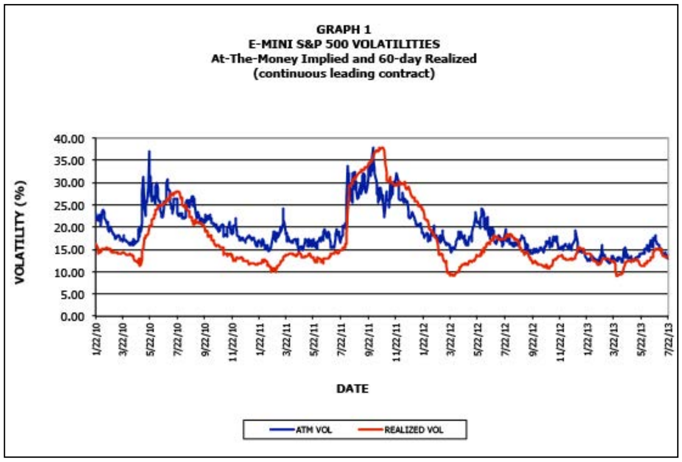

Source: Bloomberg

Note: Implied volatility is interpolated by Bloomberg as a constant 3-month time to expiration, so the volatilities shown are not actual implied volatilities of any particular option.

Graph 1 illustrates the movement of both realized and implied volatility in recent years in the E-mini S&P 500 futures market. (Implied volatility here is constant 3-month at-the money – or ATM – implied volatility.) They clearly move similarly. Despite this, the option market's implied volatility is not a predictor for historical/realized volatility in the near future. As Graph 1 makes clear, often one measure of volatility will rise while the other falls without obvious predictive power.

III. BASIC VOLATILITY TRADING

Before looking at specific types of volatility transactions, we should review a few matters that will figure in the analysis to come.

Implied volatility often rises noticeably as we approach major economic reports and policy announcements – for instance, the monthly Unemployment Rate and Nonfarm Payroll Employment reports and policy announcements by the Federal Reserve Open Market Committee. Since these reports can profoundly change market opinion about the strength of the U.S. economy they can have an explosive influence on equity prices, and therefore on the S&P 500 index and other indexes, on bond yields, and on foreign exchange levels. Traders will often buy volatility as these events approach, even early on the morning of the release, expecting – hoping, one may say-- that the news will be outside of expectations, produce a sharp move on the index level, and reward those who bought the volatility. After the news is digested the implied volatility normally declines again either because the news turned out to be not so important, or it was important and those who had bought volatility liquidate their positions. This is one of the most important and frequent occasions for short-term volatility trading not only in E-mini S&P 500 options but in many other option markets.

Graph 2 illustrates a particularly dramatic experience of this sort. In August 2011 when a Congressional impasse over an increase in the U.S. Federal debt limit and the subsequent downgrade in the credit rating of U.S. Treasury securities by Standard & Poor caused both implied and realized volatility for E-mini S&P options to spike upward by about 10 percentage points, from about 18% to more than 30%, and several months later to return to the earlier level.

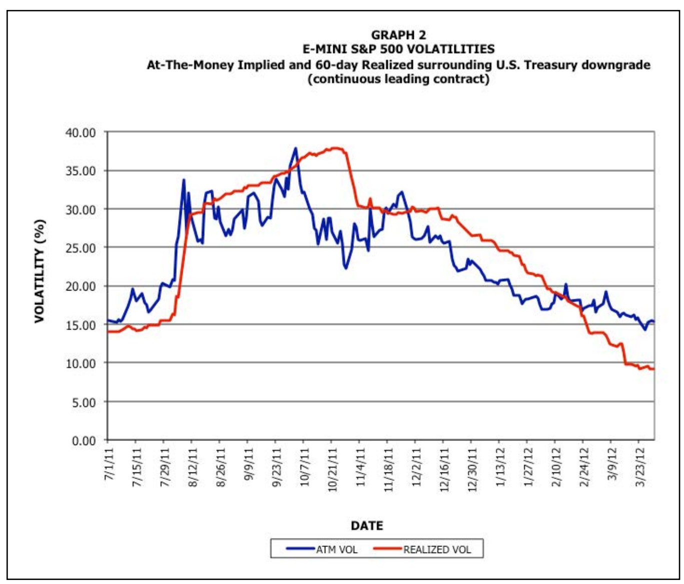

Source: Bloomberg

Note: Implied volatility is interpolated by Bloomberg as a constant 3-month time to expiration, so the volatilities shown are not actual implied volatilities of any particular option.

As an example of how great the impact of a rapid change in volatility can be, look at Graph 2. If one anticipated an increase in volatility in late July 2011 as arguments over the U.S. debt limit intensified and bought an option "straddle" to trade on that assumption what would have been the result? (A straddle is explained below, but I imagine that most readers are familiar with the principle.) In this case we'll consider buying the August 2011 1335 call and 1335 put on 25 July 2011 and liquidating the trade on 8 August 2011. As we now know from that graph, volatility – both implied and realized -- jumped in the next couple of weeks from about 18% to more than 30%. So the results would have been:

TABLE 1				
E-MINI S&P 500 STRADDLE PURCHASE				
		7/25/11	**Amount**	
Bought 1 August 2011 1335 Call		23.50	$1,175.00	(= 23.50 x $50.00)
Bought 1 August 2011 1335 Put		25.00	**$1,250.00**	(= 25.00 x $50.00)
ESM3 futures price: 1333.50				
	Total	48.50	**$2,425.00**	
		8/8/11		
Sold 1 August 2011 1335 Call		0.00	$0.00	(=0.00 x $50.00)
Sold 1 August 2011 1335 Put		223.70	**$11,185.00**	(=223.70 x $50.00)
ESM3 futures price: 1111.25				
	Total	**223.70**	**$11,185.00**	
Profit (Loss)			**$8,760.00**	

Keep in mind that this was an abnormally large jump in volatility in a short time in response to a major economic development. But this illustration is a dramatic example of how potent a rapid change in volatility can be.

Second, the *sensitivity* of an option premium to volatility will be important. In general options with a long time to expiration are more sensitive than options with only a short time left. This suggests that volatility trades are often more attractive when placed in distant expiration months. Also the sensitivity is greatest for at-the-money (ATM) options, far less so for deep in-the-money (ITM) and deep out-of-the-money (OTM) options. The measure for sensitivity to volatility is *vega*, one of the option "greeks" (although vega is not a Greek letter, or even a Greek word – it's the name derived from Arabic of an important star), the partial derivative of option premium with respect to volatility ($\partial \mathrm{Pr}/\partial\sigma$). Graph 3 illustrates the strength of vega for ATM options compared to ITM and OTM strikes, and compares its strength for 3-month and 6-month expirations. Also note that the graph is almost symmetrical for strike prices higher and lower than at-the-money.

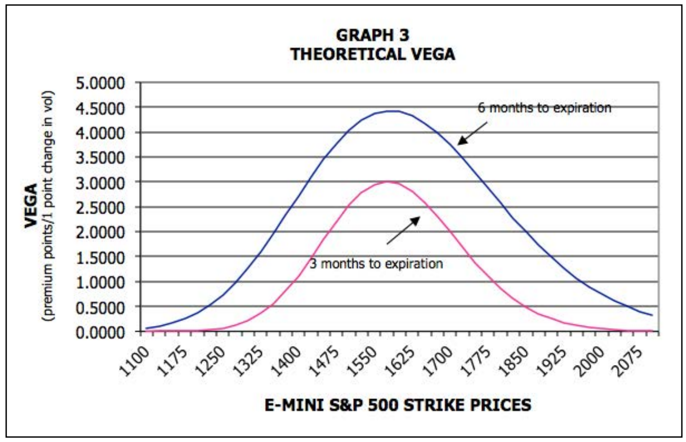

Source: author's calculations

IV. WHAT WILL VOLATILITY BE?

Volatility traders want to know whether implied volatility will rise or fall. To make such predictions – always guesswork – they often use techniques similar to those which traders employ to predict price movements. For example, they may analyze the fundamentals of the economy or fiscal and monetary policy (Will a Federal Reserve policy become less predictable and cause greater volatility? Will a confident uptrend in stock prices calm markets and reduce volatility?)

Technical analysts can apply the many tools of technical price analysis to volatility as well, treating the time series of volatility in the same way as the time series of prices. In Graph 1, an analyst might see major technical "support" for volatility at around 15% (which was valid for several months and was followed by a spike higher), or later at 13%. Or one can see a trendline traced on a graph as in Graph 4, and conclude that if it is broken (as it was) that is a strong suggestion that volatility will decline further (as it did).

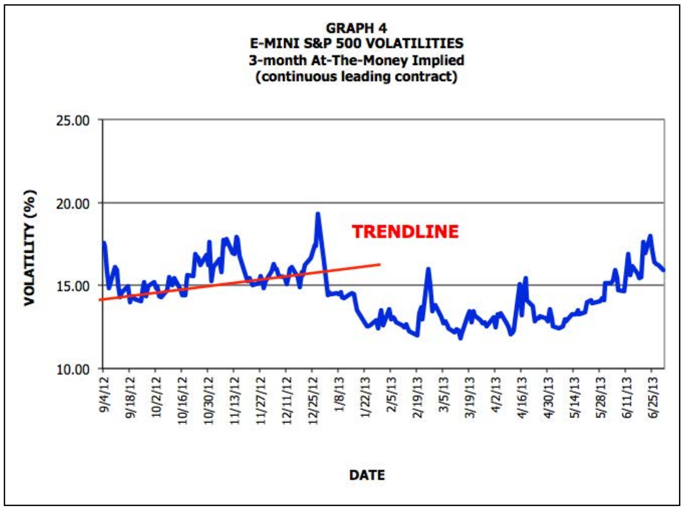

Source: Bloomberg data

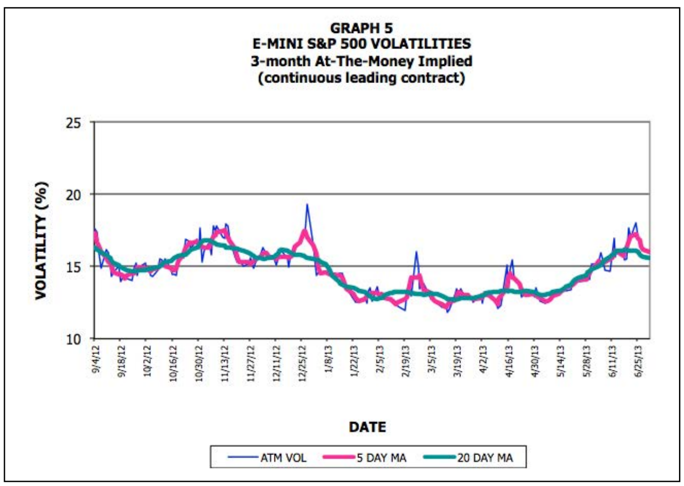

Source: Bloomberg data, author's calculations

Or one can follow moving averages of volatility, as in Graph 5, and use their smoothing effect to indicate a trend while ignoring some of the day-to-day fluctuations or "noise". As in price analysis, any such technique is approximate and only partially scientific, but many traders find such approaches useful.

One other tool, however, is unique to examining volatility: "volatility cones". The "cone" represents historical volatility for the E-mini S&P 500 market for various spans (in Graph 6, for 20 trading days – roughly one calendar month – 40 days, 60 days, etc. Readers should note that the horizontal axis shows days until expiration declining to the left.) showing the average volatility at each time span as well as the corresponding maximum and minimum.

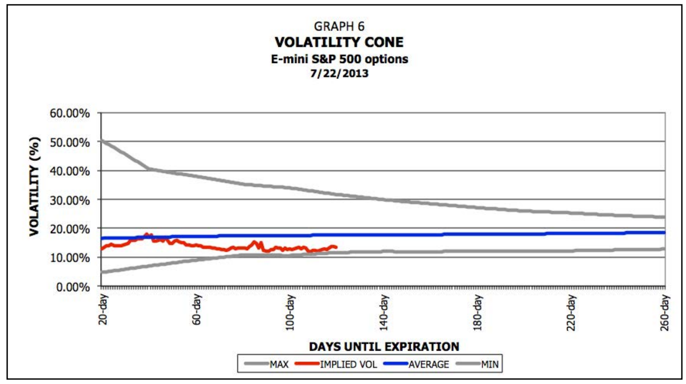

Source of data: Bloomberg. Author's calculations

These comprise the "cone" itself, showing approximate ranges within which realized volatility tends to move. Superimposed on that cone is the implied volatility of the same market. If the implied volatility is near the maximum – the top of the cone – a trader will reasonably infer that IV is likely to decline. Then any option strategy the trader designs can incorporate the probability of declining volatility. The strategy may be designed to exploit falling implied volatility, which would enhance performance (see Burghardt and Lane, 1990). At the least, the trader will see the risk that IV will decline and impair performance. Conversely, if IV is near the bottom of the cone, the trader might construct a strategy that would benefit from increasing volatility. As Graph 6 shows, recent implied volatility – this is from late July 2013 – has been near the bottom of the cone and most recently is near the average. Consequently as of this writing the volatility cone does not indicate much direction for implied volatility for the near future.

Traders might construct any of these tools themselves or rely on brokers, advisors, etc., to provide them.

V. SIMPLE STRATEGIES

Readers of this note are surely familiar with the basic method of "buying" implied volatility when it is low (compared to its recent performance) or "selling" it when it is high, expecting that it will return to something like its recent average, through straddles and strangles. Let us review these concepts, briefly:

The trader first has to judge whether implied volatility will stay at the current level (about 16-17% at the time of writing), increase, or decline. There will be a million ways to approach that, as suggested in the previous section. The trader, though, has to look at this question in regard to the period for which he will hold a position (the fact that IV reached 60-70% in 2009 is probably not relevant for a short-term or intermediate-term trade today, but the fact that it was a very low 11% in May 2013 may be relevant).

One buys or sells a *straddle* by buying or selling both a call and a put with the same strike price in one expiration month. The profit/loss profiles of the straddle are familiar (Graph 7). The buyer pays the premium for both the put and the call and loses that amount if the underlying futures price does not move much before expiration, but makes a profit if the underlying either rises or falls enough – outside the "breakeven" prices – to compensate for that cost. This is one way in which the straddle buyer "buys volatility" – he wants the underlying price to move but doesn't care whether it rises or falls. The straddle seller receives both of those option premiums, so he has "sold volatility" since his trade relies on the prospect that underlying futures price fluctuates little.

GRAPH 7

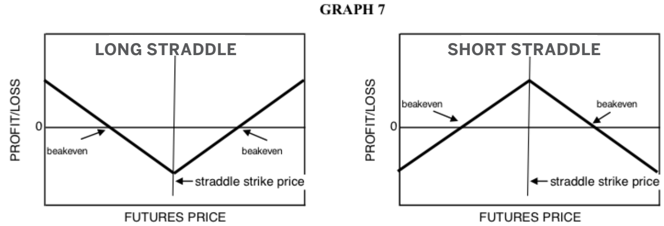

The *strangle* is identical in principle but here the strike prices are different: the strangle buyer buys a call with a higher strike price and a put with a lower strike price. The strangle seller sells those options. The profit/loss profiles (Graph 8) are like those for the straddle except that the middle portion is a flat area between the strike prices and the breakeven prices are farther apart than in the straddle.

GRAPH 8

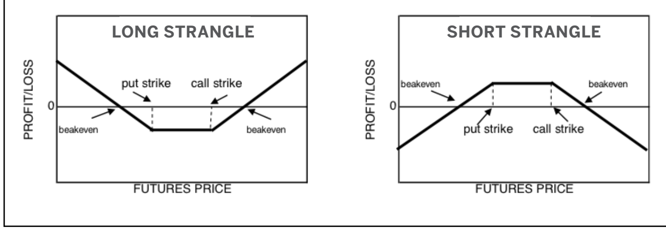

In both strategies the buyer wants prices to move a lot and in this sense he buys volatility, and in the same sense the seller sells volatility. However, implied volatility – not merely actual price movement – is also a factor in the straddle/strangle strategy. Let's first look at the seller.

Any change in implied volatility makes no difference at all to the profit/loss profiles *at expiration* – implied volatility is irrelevant then. *Before expiration*, however, a change in implied volatility will change the profit/loss profile at all underlying futures prices. IV, remember, can increase even if realized HV hasn't changed (for instance, if market participants, expecting the upcoming Fed meeting to change U.S. monetary policy, continue to buy volatility in anticipation of that).

Graph 9 illustrates the profit/loss profile of a short straddle in September 2013 E-mini S&P 500 options (ESU3) from 1 July 2013 using the ATM 1615 put and call with volatility of 15.64%. The blue curve shows profit and loss estimated as of two months later on 1 September with no change in IV; the green curve with a 2 percentage point decline in IV to 13.64%, the dashed curve at expiration. Over that time the decline in implied volatility increases profit at every level of the underlying price, combining with time decay to work in favor of the position.

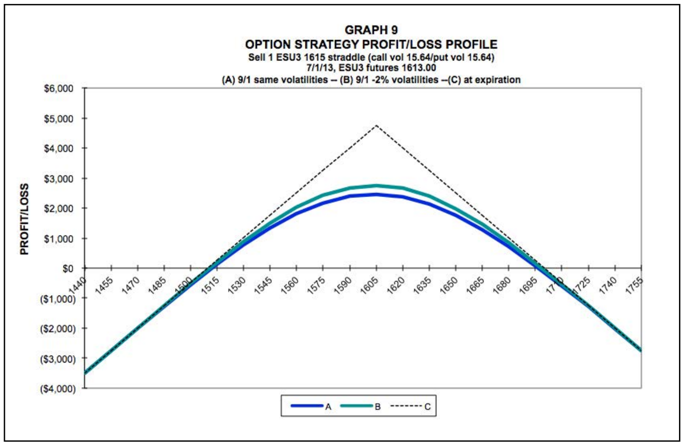

Now let's consider the buyer of this same straddle. Graph 10 represents profit/loss on two dates – one just two days after the transaction, the other on 1 September. The blue curve here shows P/L on 3 July with no change in volatility; the green curve represents P/L on that date with a 2 percentage point increase in IV to 17.64% -- profit is greater despite the fact that there have been two days of time decay. But the red curve is the interesting one: it shows P/L on 1 September *with the same 2 point increase in volatility*. Profit is much lower at every point – the big increase in IV has been overwhelmed by time decay. Where the profit/loss profile shows profit for this date it is because of the move in prices, perhaps an increase in realized, historical volatility, not because of the increase in implied volatility. To be sure, the profit is higher for that date than if IV had not increased (that curve isn't shown in the graph), but not enough to overcome the loss from time decay.

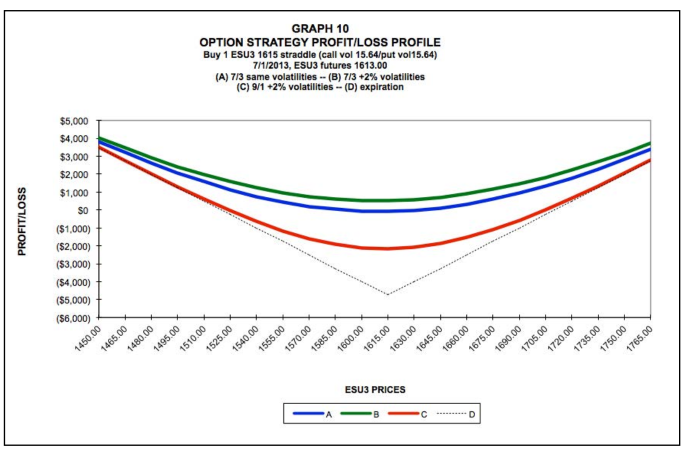

So if the buyer wants IV to increase he should hope for it to occur soon – option prices will respond substantially to an increase in IV when there is a lot of time left in the life of the option, but far less later on, as is clear from Graph 2 of vega – vega is much lower at all prices for the shorter-term option.

The upshot of this explanation is that if you foresee volatility declining you can take a position and be patient; if you foresee it increasing it is important to take a position not long before that happens.

VI. SKEW

When we say that an option's implied volatility is a form of market estimate of the realized volatility of the E-mini S&P 500 price between now and the option expiration, it would be reasonable to anticipate that the estimate would be the same for every option of a given expiration, for every strike price, both puts and calls. In fact, it is *not* the same for all those options. Furthermore, the differences among implied volatilities are not trivial -- they are far greater than one might attribute to bid/ask spreads or transitory market flows. They are persistent and more-or-less systematic. Graph 11 shows these implied volatilities for E-mini S&P 500 options on a recent day:

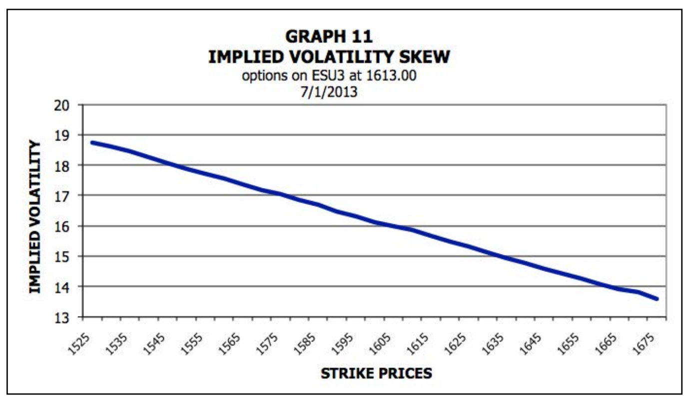

Source: QuikStrike

It's clear that the implied volatilities differ a lot from one strike price to another; also that the volatility for low-strike-price options is much higher than the volatility of at-the-money options, which in turn are higher than volatilities for high-strike-price options. This is called the "volatility skew" and is typical for options on stocks, stock indexes, and stock index futures.* (For some markets the volatility skew may go in the opposite direction, with volatilities rising as strike prices rise – this appears frequently for commodity options; for still others the implied volatilities may be higher than at-the-money volatilities on both sides of the at-the-money strike price – that formation is usually called a “volatility smile”.)

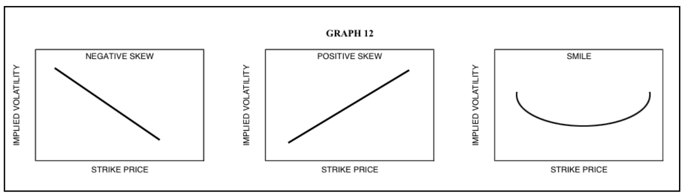

Several explanations for this anomaly have been put forward – and an “anomaly” it is: option pricing theory is based in part on a supposition that prices evolve over time according to a stable stochastic process with a fixed volatility – but none of the explanations seem satisfying. There are models of “stochastic volatility”, treating the volatility itself as random. The skew may be a manifestation of market flows as hedgers buy out-of-the-money puts to protect stock portfolios. This explanation has some appeal since index volatilities showed a horizontal line or even a “smile” before the shock of the 1987 stock market crash and have shown a negative skew since then, suggesting persistent fear of similar crashes. Related to this, it may reflect the fact that actual price behavior embodies “fatter tails” than the lognormal distribution typically presumed in option theory. But the source of the anomaly is not central to our purpose here – we want to examine how an investor or trader can make use of this phenomenon.

The volatility skew can intensify relative value strategies – vertical call or put spreads, ratio spreads, backspreads, etc. “Intensify” is the key word. The trader must first believe in the basic structure of the strategy, before considering the influence of the skew. Take a look at one simple and common combination: a “bear call spread”. (Don’t think about whether this position is attractive at this moment; we only want to see how the volatility skew can affect it.)

On 1 July, when the E-mini S&P futures price settles at 1606.75, a bearish trader wants a limited-risk bearish option position and so . . .

Sells 1 ESU3 1535 call (IV = 18.40%)

Buys 1 ESU3 1620 call (IV = 15.43%)

The basic profit/loss profile for this structure is

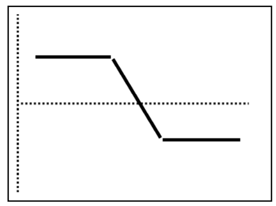

. . . it is bearish, but has less risk from a price rise than a naked short call.

The position sells a low-strike-price, high-volatility call and buys a high-strike-price, low-volatility call, so its results can be affected by the skew between the volatilities. When we evaluate this (theoretically) – shown in Graph 13 — after two months for 1 September, we see the blue curve for the case in which the implied volatilities of the two legs do not change, a higher green curve if the volatility skew narrows by 2 percentage points (to 17.40% and 16.45%), and lower red curve if the skew widens by 2 percentage points (to 19.40% and 14.45%). The dashed curve is the profile at expiration.

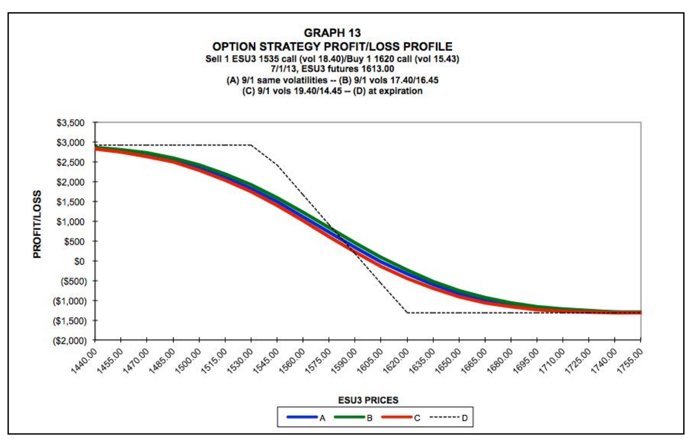

Do not confuse this with the statistical concept of “skewness” – they are not related.

Three points about this exercise are noteworthy: (1) the movement of the skew improves profit (if skew contracts) or reduces it (if skew worsens); (2) the change in profit is not great, about $100-to-$150 at the most either better or worse; (3) this small difference in profitability is from a full 2 percentage point change in skew, which would be a pretty big change.

That is why I referred to volatility skew as "intensifying" a transaction – the position will be a winner or a loser based primarily on the movement in the underlying ESU3 price and the structure of the position. The skew affects the outcome but only modestly. If we were to duplicate this exercise for other structures involving the skew we would see the same pattern of a helpful but small effect. As we think back to the graph of option vegas, this result is reasonable: in structures of this sort we would generally be using ITM and OTM options where vega is relatively low compared to ATM vega, so the impact of a change in volatility, even a quite large change, will not significantly alter the trade's results.

Furthermore, keep in mind that traders who are adept at using the volatility skew treat the difference in volatilities like other market spreads: sell the relatively high volatility and buy the relatively low volatility. To do this they don't merely look at the absolute levels of volatilities – they won't naturally converge to one level. They want to judge whether today's volatility difference, or skew, or volatility curve slope is unusually wide or unusually narrow. For this they have to construct histories of the skew that cover many option expirations and a wide range of outright index levels and exercise prices. The skew may have been steepening and will continue to steepen – present levels tell us nothing about that. The volatility traders will try to see whatever pattern there may be in a history of the volatility curve.

To do this it is necessary to standardize the volatility curve across expirations and price levels. There are many methods for doing that but the overall principle is something like this:

- Compare volatilities for options whose strike prices are equally far away from the at-the-money strike, perhaps 5% lower and 5% higher. This is often referred to as "95% moneyness" and "105% moneyness". They may look at the outright difference in volatility, or divide that difference by the volatility of the ATM options to obtain a kind of ratio or index.
- Compare volatilities of the put whose delta is -0.25 (an out-of-the-money put; its strike price is far below that of the at-the-money option) and the volatility of the call whose delta is +0.25 (also out-of-the-money; its strike price is far higher than at-the-money). This also may be viewed as the outright difference of volatilities or divided by the volatility of the at-the-money options. Each volatility trader has a favorite measurement of skew but all aim to see the slope of the volatility curve in a way that can be compared over time.

Graph 14 shows how this kind of measure has moved in the last couple of years (an index of the volatility for 95% moneyness minus 105% moneyness divided by at-the-money volatility).

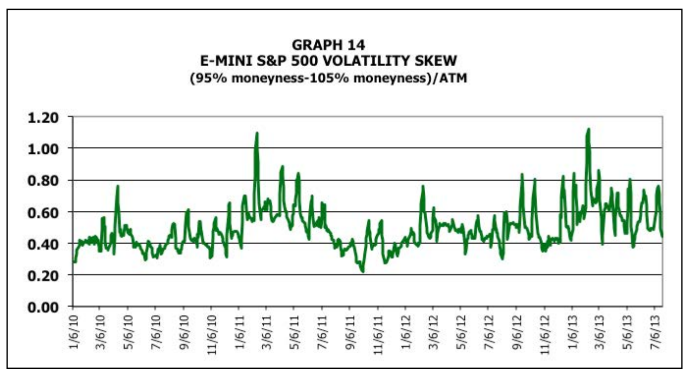

Source: author's calculations based on Bloomberg option premium data.

Obviously this shows a great deal of variation – in fact, the graph shows a 5-day moving average of that index because the raw numbers are so choppy. The point here is that from examining a history like this a trader can see when the skew is unusually great or unusually narrow. For instance, you may want to construct a position that will benefit from a narrowing of this "spread" when it is at 0.70 or higher; or structure a trade that benefits from an increasing skew when it is near

0.30; or other tools of technical analysis may come into play – trend lines, moving averages, etc. – to find opportunities where movement in the skew can be expected to “intensify” the profitability or reliability of a relative value position that is attractive on its own terms, as mentioned earlier.

One additional and very interesting inference to draw from this graph is that the level of “skewness” reverses quickly from extreme to more central levels, suggesting that use of the skew in trade design is normally a short-term matter.

VII. CONCLUSION

The principal conclusions to draw from this brief survey of option volatility in the E-mini S&P 500 market are that . . .

1. Volatility, both implied and realized, is an important element that affects option pricing and must always be taken into account in using options as a speculative vehicle or a hedging vehicle, or in relative-value positions.

2. Implied volatility is most powerful (i.e. vega is greatest) for at-the-money options with relatively long times to expiration.

3. Strategies which rely on buying volatility should be designed to exploit volatility increases in the short-term, while those selling volatility can be more long-term.

4. Exploiting volatility skew can intensify the results of option structures but does not provide the major part of performance.

5. In attempting to predict movements in volatility, traders can apply many of the same techniques used to forecast price movements. However, volatility cones are a tool unique to options analysis.

REFERENCE

Burghardt, Galen and Morton Lane (1990), “How to tell if options are cheap”, *Journal of Portfolio Management*, Volume 16, Number 2 (Winter),pp. 72-78.

Lawrence Morgan is Adjunct Instructor in Finance at DePaul University and in Economics at Dominican University, following a 33-year career in the financial futures and options brokerage industry assisting clients in devising and executing hedging and trading strategies, primarily in fixed-income and currency markets. His clientele included commercial and investment banks, hedge funds, insurance firms, and corporate treasurers in North and South America, Europe, and East Asia.

He has incorporated derivatives analysis and training throughout his career, instructing clients and others in the theory and practice of derivatives markets, as part of the instructional staff at Dean Witter Reynolds as well as in courses through the Korea Banking Institute in Seoul and at the People’s University of Beijing. Before entering the brokerage profession, he was Senior Staff Economist at the Chicago Board of Trade.