GLOBAL VOLATILITY SUMMIT 2013

April 2013 Newsletter

2014 Event Details

Date. Plans are already underway to secure a convenient date for the 2014 event!

Please continue to check the website for registration, updates and tentative agenda (www.globalvolatilitysummit.com).

2013 Event Recap

Keynote speaker. Sal Khan, founder of The Khan Academy and author of *The One World Schoolhouse* gave an insightful presentation on using technology to innovate the way education is provided across the globe.

Special Guest Speaker. Mike Edleson followed up to his 2012 GVS talk about the decision to implement a tail hedge, with an informative discussion on implementation of a tail hedge and how to identify the right managers for your mandate. Mr. Edleson's presentation is available on the GVS website.

Managers. The following managers participated:

Blue Mountain Capital
Capstone Investment Advisors
Fortress Investment Group
Forty4 Fund
Ionic Capital Management
JD Capital Management
Parallax Fund
PIMCO
Pine River Capital Management
Saiers Capital

Questions?
Please contact info@globalvolatilitysummit.com

2013 Event Summary and April research piece

The fourth annual Global Volatility Summit ("GVS") was a success. The event took place on February 25th in New York City, and ten volatility and tail hedge managers hosted an audience of over 350 people. The event featured a thought provoking key note speech by Sal Khan regarding the transformation of the educational process to a web based mode of communication, a presentation by Mike Edleson from The University of Chicago on tail hedging implementation, and four panels including a pension and consultant panel.

The primary goal of the GVS is to educate the investment community about volatility and how it can help investors attain their growth targets. The GVS is an evolving community of managers, investors, and industry experts. We rely on the feedback and guidance of our investors to shape the event and line-up of speakers each year. Following the summit in February, a number of you requested more fundamental knowledge on volatility trading strategies. As a result, we are sharing a comprehensive piece on volatility trading strategies co-authored by Colin Bennett and Miguel Gil of Santander. We thank them for sharing this piece, which we believe you will find to be informative.

If you have any topics you would like to see the managers address in future newsletters please send us an email.

Cheers,
Global Volatility Summit

Santander GLOBAL BANKING & MARKETS

Equity Derivatives **Europe**

Madrid, November 5, 2012

VOLATILITY TRADING

Trading Volatility, Correlation, Term Structure and Skew

Colin Bennett
Head of Derivatives Strategy
(+34) 91 289 3056
cdbennett@gruposantander.com

Miguel A. Gil
Equity Derivatives Strategy
(+34) 91 289 5515
mgil@gruposantander.com

VOLATILITY FUTURE ETN/ETF

Structured products based on constant maturity volatility futures have become increasingly popular and in the US have at times had a greater size than the underlying volatility futures market. As a constant maturity volatility product needs to sell near-dated expiries and buy far-dated expiries, this flow supports term structure for volatility futures and the underlying options on the index itself. The success of VIX-based products has led to their size being approximately two-thirds of the vega of the relevant VIX futures market (which is a similar size to the net listed S&P500 market) and, hence, appears to be artificially lifting near-dated term structure. The size of vStoxx products is not yet sufficient to significantly impact the market, hence they are a more viable method of trading volatility in our view. We recommend shorting VIX-based structured products to profit from this imbalance, potentially against long vStoxx based products as a hedge. Investors who wish to be long VIX futures should consider the front-month and fourth-month maturities, as their values are likely to be depressed from structured flow.

STRUCTURED PRODUCTS ON VOL FUTURES IMPROVED LIQUIDITY

Size of VIX-based products has at times been greater than underlying VIX futures market

As it is impossible to have a product (perpetual or otherwise) whose payout is the volatility index itself, volatility futures were launched to give investors an easy method of trading volatility. Initially, VIX and vStoxx futures had limited liquidity, potentially as they are not perpetual; however, the creation of perpetual structured products has improved the liquidity of volatility futures. Similarly, the introduction of options on these futures has increased the need to delta hedge using these futures, also increasing liquidity. In the US, the size of structured products on VIX futures is so large at times it was bigger than the underlying VIX futures market and appears to have moved the underlying S&P500 market itself.

VIX PRODUCTS ACCOUNT FOR 2/3 OF THE SIZE OF FUTURES MARKET

The size in vega of the US market for vanilla S&P500 options, VIX futures and VIX-based ETN/ETF is shown in Figure 69 below. As can be seen, the size of VIX-based ETN/ETFs is approximately two-thirds of the size of the relevant VIX future.

Figure 69. S&P500 Vega by Maturity for Options, Volatility Futures and ETN/ETFs

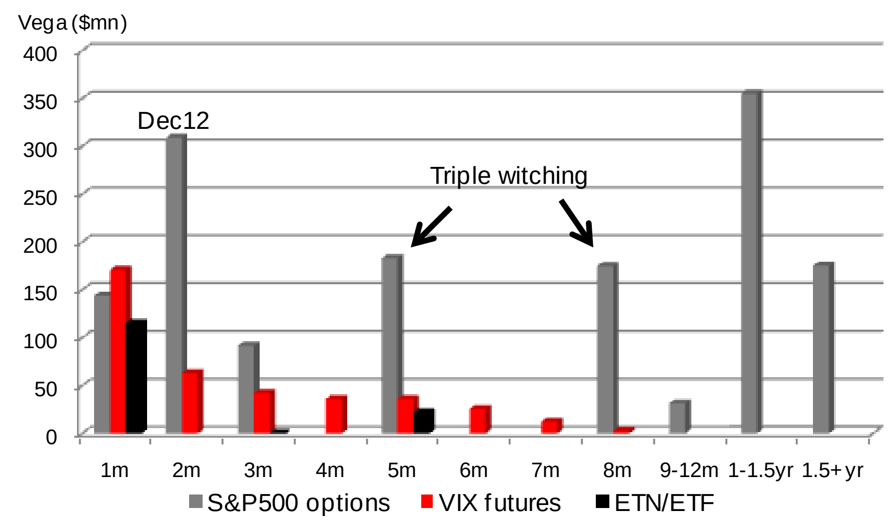

Source: Santander Investment Bolsa.

VIX FUTURES MARKET IS NOW SIMILAR SIZE TO NET OPTIONS MARKET

While the size in vega of the S&P500 options market is c4 times bigger than the VIX futures market, this is a theoretical maximum for the market. In practice, if investors trade a spread (eg, call/put spread or ladder) the vega of these structures is far less than the combined vega of the individual legs (vega is the difference between the long and short legs, not the sum of the legs). There is also significant trade in synthetics (long call short put as a substitute for long future) for non-triple witching expiries as S&P500 futures only exist for quarterly expiries. In addition, as one cannot cross futures, when trading on swap (trade volatility structure delta hedged so price is not affected by movements in spot) the delta hedge is done via synthetics. Interest rate trades such as box spreads (long synthetic of one maturity and short synthetic of another maturity) also have no volatility component. A reasonable assumption is that the size of the net vega of the S&P500 listed options market is c25% of the theoretical maximum. Hence, the size of the net listed S&P500 vega is similar to that of the VIX futures market.

VIX futures size compared to S&P500 reduces if OTC market is taken into account

However, we estimate that the OTC market for the S&P500 is 50%-100% of the size of the listed market. This is due to the significant long-term hedging (eg, from variable annuity programs) which cannot be done on exchange (as only maturities up to 2-3 years are listed on the S&P500). Additionally, the size of the variance swap market adds to the size of the OTC market. Hence, we estimate the vega of VIX futures would be 50%-100% of the total (listed and OTC) size of the S&P500 market.

Figure 70. Volatility Future Term Structure

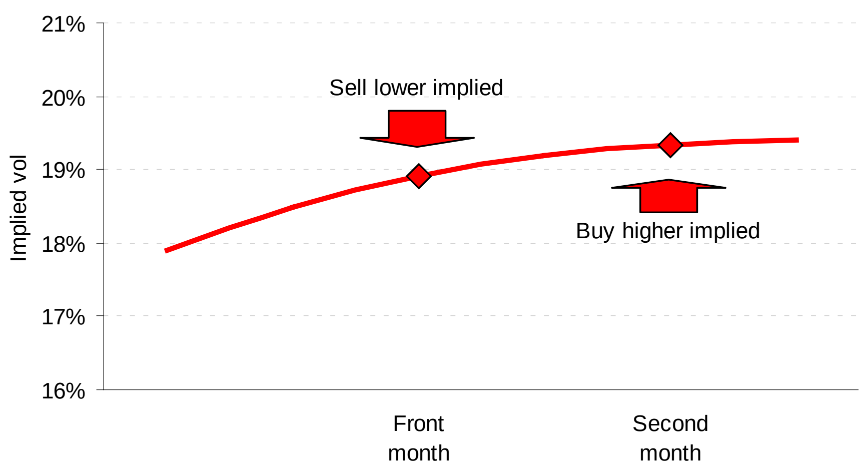

Source: Santander Investment Bolsa.

OPEN-ENDED VOLATILITY PRODUCTS STEEPEN TERM STRUCTURE

Structured products often sell near-dated expiries and buy longer-dated expiries, lifting term structure

While futures on a volatility index have the advantage of being a listed instrument, they have the disadvantage of having an expiry and, therefore, a longer-term position needs to be rolled. In response to investor demand, many investment banks sold products based on having a fixed maturity exposure on an underlying volatility index. As time passes, these banks hedge their exposure by selling a near-dated expiry and buying a far-dated expiry. The weighted average maturity is therefore kept constant, but the flow puts upward pressure on the term structure. For products of sufficient size, the impact of structured products on the market ensures the market moves against them. Products on short-dated VIX futures which have an average 1-month maturity (by selling front month and buying the second expiry) are now sufficiently large to be moving the volatility market for the S&P500.

Figure 71. Size of VIX ETN/ETF, 2009-12

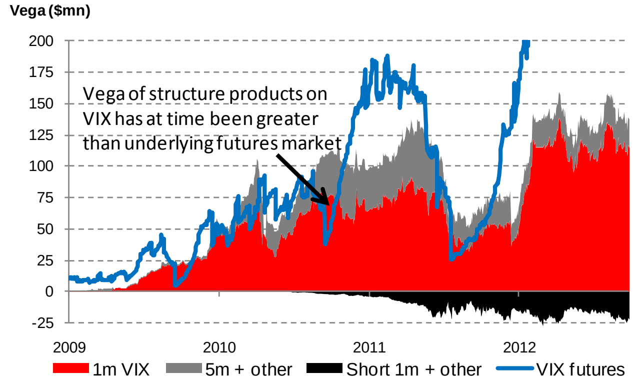

Size of vStoxx ETN/ETF, Mar10-Sep12

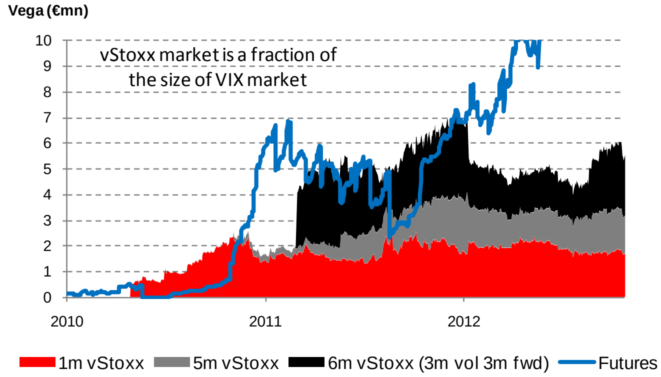

Source: Santander Investment Bolsa estimates.

Volatility products suffer 'roll-down cost' due to positive term structure

As vStoxx based ETN/ETF are not as large as VIX products, they do not suffer from the same supply/demand imbalance

Historically, due to risk aversion and supply demand imbalances, the average term structure of implied volatility has been positive. Index implied volatility term structure will also be lifted by positive implied correlation term structure. The launch of volatility futures and their related ETN/ETF products has increased the supply demand imbalance and supported term structure for the S&P500. VStoxx products are not yet large enough to have an impact on SX5E term structure. As long volatility ETN/ETF products are always selling near-dated implied and buying far-dated implied, there is a roll-down cost if the term structure is positive (as a low near-dated volatility future is sold and a high far-dated volatility future is bought). The higher the positive term structure, the greater the roll-down cost. Conversely, volatility ETN/ETF products will benefit from negative term structure. Investors tend not to benefit from the periods of time there is positive roll-down cost, as these products are often used as a hedge (or view on volatility increasing) and the position is typically closed if equity markets decline and volatility spikes.

Figure 72. SX5E Vega by Maturity for Options, Volatility Futures and ETN/ETF

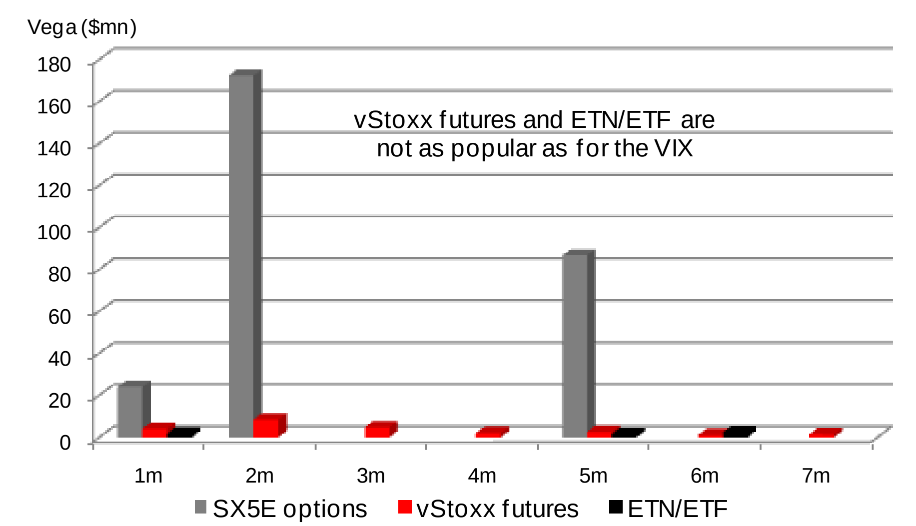

Source: Santander Investment Bolsa.

Longer dated maturity ETN/ETF products have less roll-down cost, but lower beta

Given the fact term structure flattens out as maturity increases, the gradient of volatility surfaces is steeper for near-term expiries than far-term expiries. This means structured products based on longer-dated volatility suffer less roll-down cost. This led to the creation of medium-dated ETN/ETFs, which have an average maturity of five months as they invest in futures of maturity 4, 5, 6 and 7 month (hence sells 4th future to buy 7th). However, as the beta of volatility futures decreases with maturity, longer-dated volatility products benefit less from volatility spikes. While the ratio of beta to roll-down cost is similar across different maturity volatility products, near-dated products do have a worse ratio. There are some products that try to benefit from the excess demand for near-term 1-month ETN/ETFs (medium-dated 5-month ETN/ETFs are less popular) by going short a 1-month volatility product, and at the same time going long approximately twice that size of a 5-month volatility product (as $\sqrt{5}\approx 2$ and volatility often moves in a square root of time manner).

Figure 73. Flows of Volatility Based ETN/ETF

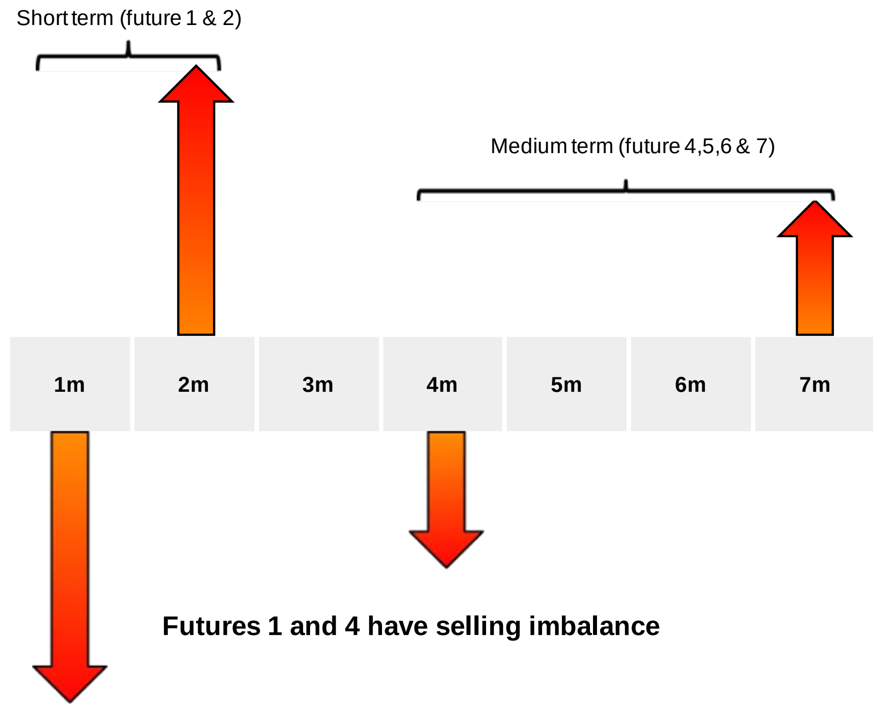

Source: Santander Investment Bolsa.

BUYING FRONT MONTH OR 4TH MONTH VOLATILITY FUTURES IS BEST

Long fourth volatility future should be closed when it becomes second expiry

If an investor wishes to initiate a long volatility future position, the best future is the front-month future as this has the most selling pressure from ETN/ETFs (and hence is likely to be relatively cheap). For investors who wish to have a longer-dated exposure, we would recommend the fourth volatility future as our second favourite. This future benefits from the selling of medium-dated ETN/ETFs. A long fourth future position should be closed when it becomes the second month future (as this price is supported by the short dated ETN/ETF). While the eight future is also a viable investment, the liquidity at this maturity is lower than the others.

WE RECOMMEND SHORTING NEAR-DATED VIX BASED ETN/ETFS

Short VIX long vStoxx profits from VIX imbalances (while hedging overall volatility levels)

Given the imbalance in the VIX futures market resulting from the size of ETN/ETF products for near-dated VIX products, we recommend shorting these products. The XXV ETF (inverse of the VXX, whose ticker is also the letters of the VXX backwards) based on short 1-month VIX futures is also a viable method of profiting from this imbalance. As the size of the XXV is only c20% of the size of the VXX, a significant imbalance still remains in our view. As vStoxx-based products are not sufficiently large to be causing an imbalance, a short VIX product long vStoxx product is an attractive way to profit from the VIX imbalance while hedging the overall level of volatility (we note that this trade does not hedge any US or Europe specific volatility). As can be seen in Figure 74 below, the profile of 1m vStoxx/1m VIX (proxy for long 1m vStoxx and short 1m VIX rebalanced every day) offers an attractive performance.

Figure 74. VIX and vStoxx 1-Month Rolling Volatility Future Performance (Rebased)

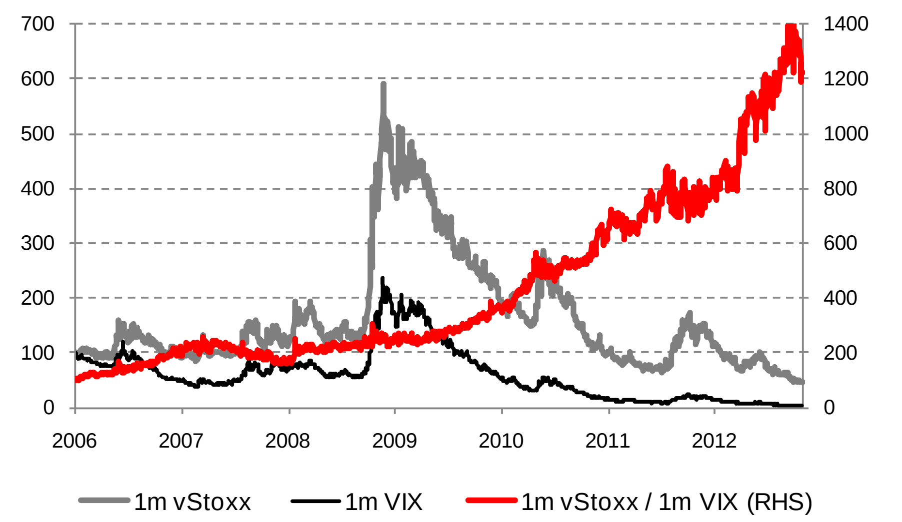

Source: Santander Investment Bolsa.

ETN HAS COUNTERPARTY RISK UNLIKE ETF

There are both ETNs and ETFs based on volatility futures, the primary difference being counterparty risk. Despite the fact the underlying of the product is listed (and hence has no counterparty risk), an investor in an ETN suffers the counterparty risk of the provider. ETFs do not suffer this problem.

EXCESS RETURN PRODUCTS ARE BETTER THAN TOTAL RETURN

For investors who are able to trade them, swaps based on excess return indices (eg, VST1ME for vStoxx 1-month futures) are better than ETN/ETF based on total return indices (eg, VST1MT for vStoxx 1-month futures). This is because the returns received from a total return product (EONIA) are likely to be less than the funding levels of a client.

OPTIONS ON VOLATILITY FUTURES

The arrival of options on volatility futures has encouraged trading on the underlying futures. It is important to note that an option on a volatility future is an option on future implied volatility, whereas an option on a variance swap is an option on realised volatility. As implieds always trade at a lower level to peak realised (as you never know when peak realised will occur) the volatility of implied is lower than the volatility of realised, hence options on volatility futures should trade at a lower implied than options on var. Both have significantly downward sloping term structure and positive skew. We note that the implied for options on volatility futures should not be compared to the realised of volatility indices.

OPTIONS ARE ON THE FUTURE, NOT THE VOLATILITY INDEX ITSELF

There is a significant difference in volatility between a vol index and a vol future

As volatility markets have become more liquid, investors became increasingly interested in purchasing options on volatility. As it is impossible to buy a volatility index itself, options on volatility have to be structured as an option on a volatility future. For equities there is not much difference between the volatility of spot and the volatility of a future (as futures are near dated, the effect of interest rate and implied dividend volatility is small). However, there is a very significant difference between the volatility of a vol index, and the volatility of a vol future.

VOLATILITY TERM STRUCTURE IS VERY NEGATIVE

The term structure of implied volatility of vanilla equity options is on average relatively flat[24]. In contrast, the term structure of implied volatility of option on vol futures is sharply negative. The volatility of a vol future is significantly less than the volatility of the vol index, but does converge as it approaches expiry (when it becomes as volatile as the vol index itself).

Figure 75. Realised Volatility of Vol Futures and Implied Volatility of Option on Vol Future

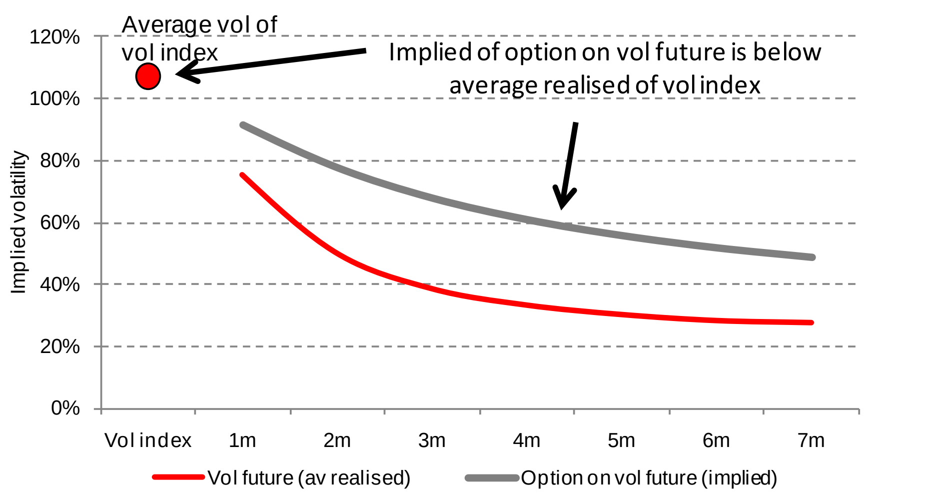

Source: Santander Investment Bolsa.

[24] On average slightly upward sloping, but at a far shallower gradient to the negative term structure of options on vol futures.

COMPARING IMPLIED VOL AND REALISED VOL IS DIFFICULT

The realised volatility of a vol future increases as it approaches expiry (near-dated volatility is more volatile than far-dated volatility). The implied volatility of an option on a vol future should trade roughly in line with the average realised volatility of the vol future over the life of the option. Hence the average realised volatility of a vol future will be a blend of the initial low realised volatility, and the higher realised volatility close to expiry. The implied volatility term structure of an option on vol future will therefore be less negative than the realised volatility term structure of the vol future.

Implied of option on vol futures is between realised of vol index and realised of vol futures

The implied volatility level of options on vol futures is also higher than the realised volatility of the vol future for that expiry (eg, implied of 6-month option on vol futures is above current realised of 6-month vol future). The implied volatility of options on vol futures will, however, be less than the realised volatility of the vol index, which makes options on vol futures look cheap if an investor mistakenly compares its implied to the realised of the vol index.

OPTION ON VOL FUTURE SHOULD BE CHEAPER THAN OPTIONS ON VAR

Implied volatility is less volatile than realised volatility, as implied volatility will never trade at the min or max level of realised (as it is an estimate of future volatility, and there is never a time that the market can be 100% certain realised will reach its min or max). As implied volatility is less volatile than realised volatility, an option on a vol future should be at a lower implied than an option on realised variance (particularly for near-dated expiries). They will, however, have a similar negative term structure.

Options on vol future have positive skew, just like options on var

When there is an equity market panic, there tends to be large negative returns for equities and a volatility spike. As the probability distribution of equity prices has a greater probability of large negative returns, it has a negative skew. Volatility, on the other hand, tends to have a larger probability of large positive returns and hence has positive skew (just like options on realised variance).

Figure 76. Probability Distribution of Options on Vol Futures

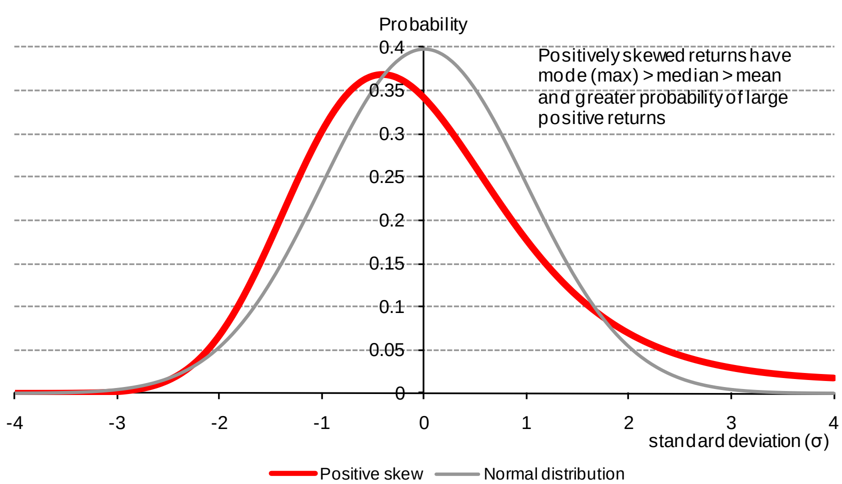

Source: Santander Investment Bolsa.

Need to model volatility with high vol of vol and mean reversion

As vol futures have a high near-term volatility, and low far-dated volatility, they have to be modelled with a high vol of vol and high mean reversion.

OPTIONS ON VOLATILITY FUTURE PRODUCTS ALSO EXIST

At present there are only options on VIX and vStoxx futures. There are, however, also options on structured products based on VIX volatility futures. The list of underlyings for options is given below.

Figure 77. Volatility Securities with Listed Options

Ticker	Underlying Type
VIX	Vol index
V2X	Vol index
VXX US	ETN
VXZ US	ETN
VIIX US	ETN
SVXY US	ETF
UVXY US	ETF
VIXM US	ETF
VIXY US	ETF

Source: Santander Investment Bolsa.

LIGHT EXOTICS

BARRIER OPTIONS

Barrier options are the most popular type of light exotic product, as they are used within structured products or to provide cheap protection. The payout of a barrier option knocks in or out depending on whether a barrier is hit. There are eight types of barrier option, but only four are commonly traded, as the remaining four have a similar price to vanilla options. Barrier puts are more popular than calls (due to structured product and protection flow), and investors like to sell visually expensive knock-in options and buy visually cheap knock-out options. Barrier options (like all light exotics) are always European (if they were American, the price would be virtually the same as a vanilla option, as the options could be exercised just before the barrier was hit).

BARRIER OPTIONS CAN HAVE DELTA OF MORE THAN ±100%

Extra hedging risk of barriers widens the bid-offer spread

The hedging of a barrier option is more involved than for vanilla options, as the delta near the barrier can be significantly more than ±100% near expiry. The extra hedging risk of barriers widens the bid-offer spread in comparison with vanilla options. Barrier options are always European and are traded OTC.

THERE ARE THREE KEY VARIABLES FOR BARRIER OPTIONS

There are three key variables to a barrier option, each of which has two possibilities. These combinations give eight types of barrier option (8=2×2×2).

- **Down/up.** The direction of the barrier in relation to spot. Almost all put barriers are down barriers and, similarly, almost all call barriers are up barriers.
- **Knock in/out.** Knock-out options have a low premium and give the impression of being cheap; hence, they are usually bought by investors. Conversely, knock-in options are visually expensive (as knock-in options are a similar price to a vanilla) and are usually sold by investors (through structured products). For puts, a knock-in is the most popular barrier (structured product selling of down and knock-in puts). However, for calls this is reversed and knock-outs are the most popular. Recent volatility has made knock-out products less popular than they once were, as many hit their barrier and became worthless.
- **Put/call.** The type of payout of the option. Put barriers are three to four times more popular than call barriers, due to the combination of selling from structured products (down and knock-in puts) and cheap protection buying (down and knock-out puts).

ONLY FOUR OF THE EIGHT TYPES OF BARRIER ARE USUALLY TRADED

The difference in price between a vanilla option and barrier option is only significant if the barrier occurs when the option has intrinsic value. If the only value of the option when the barrier knocks in/out is time value, then the pricing for the barrier option will be roughly equal to the vanilla option. Because of this, the naming convention for barrier options can be shortened to knock in (or out) followed by call/put (as puts normally have a down barrier, and calls an up barrier). The four main types of barrier option and their uses are shown below.

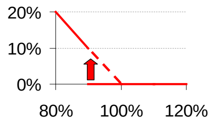

Knock-in put (down and knock-in put). Knock-in puts are the most popular type of barrier option, as autocallables are normally hedged by selling a down and knock-in put to fund the high coupon. They have a barrier which is below both spot and strike and give an identical payoff to a put only once spot has gone below the down barrier. Until spot reaches the down barrier there is no payout. However, as this area has the least intrinsic value, the theoretical price is similar to a vanilla and therefore visually expensive.

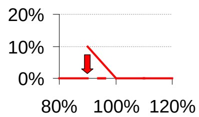

Knock-out put (down and knock-out put). Knock-out puts are the second most popular barrier option after knock-in puts (although knock-in puts are three times as popular as knock-out puts due to structured product flow). Knock-out puts give an identical payout to a put, until spot declines through the down barrier (which is below both spot and strike), in which case the knock-out option becomes worthless. As the maximum payout for a put lies below the knock-out barrier, knock-out puts are relatively cheap and are often thought of as a cheap method of gaining protection.

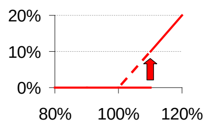

Knock-in call (up and knock-in call). Knock-in calls give an identical payout to a call, but only when spot trades above the up barrier, which lies above spot and the strike. They are the least popular barrier option, as their high price is similar to the price of a call and structured product flow is typically less keen on selling upside than downside.

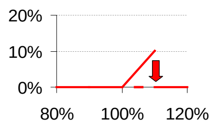

Knock-out call (up and knock-out call). Knock-out calls are the most popular barrier option for calls, but their popularity still lags behind both knock-in and knock-out puts. As they give the same upside participation as a vanilla call until the up barrier (which is above spot and strike) is reached, they can be thought of as a useful way of gaining cheap upside.

KNOCK-OUT OPTIONS DECREASE IN VALUE AS STRIKE APPROACHES

While vanilla options (and knock-in options) will increase in value as spot moves further in the money, this is not the case for knock-out options, where spot is near the strike. This effect is caused by the payout equalling zero at the barrier, which can cause delta to be of opposite sign to the vanilla option. This effect is shown below for a one-year ATM put with 80% knock-out. The peak value of the option is at c105%; hence, for values lower than that value the delta is positive not negative. This is a significant downside to using knock-out puts for protection, as their mark to market can increase (not decrease) equity sensitivity to the downside.

Figure 78. Price of One-Year ATM Put with 80% Knock-Out

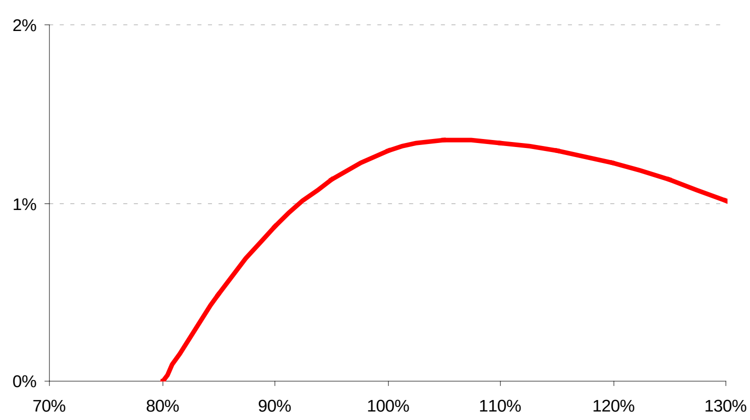

Source: Santander Investment Bolsa.

KNOCK-OUT PUT + KNOCK-IN PUT = (VANILLA) PUT

If a knock-out put and a knock-in put have the same strike and barrier, then together the combined position is equal to a long vanilla put ($P_{KO} + P_{KI} = P$). This is shown in the charts below. The same argument can apply to calls ($P_{KO} + P_{KI} = P$). This relationship allows us to see mathematically that if knock-out options are seen as visually cheap, then knock-in options must be visually expensive (as a knock-in option must be equal to the price of a vanilla less the value of a visually cheap knock-out option).

Figure 79. Knock-out Put **Knock-in Put**

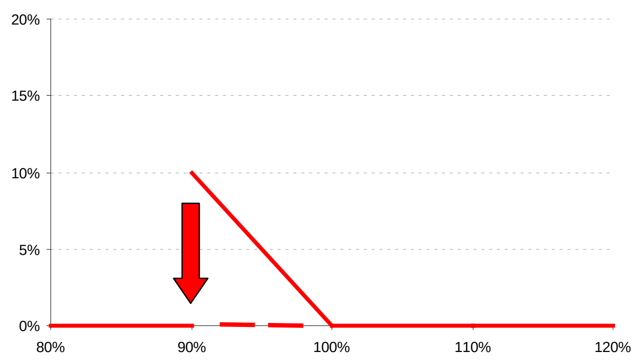

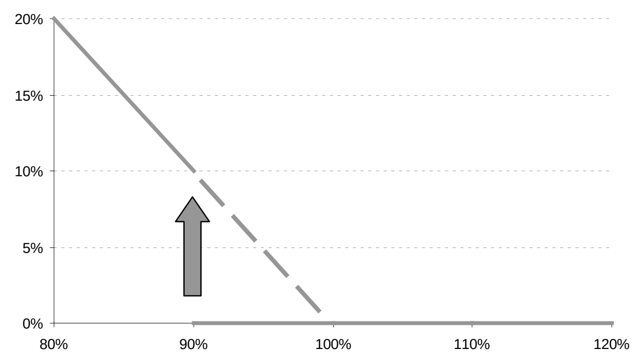

Source: Santander Investment Bolsa.

Figure 80. Knock-out Put + Knock-in Put = Put

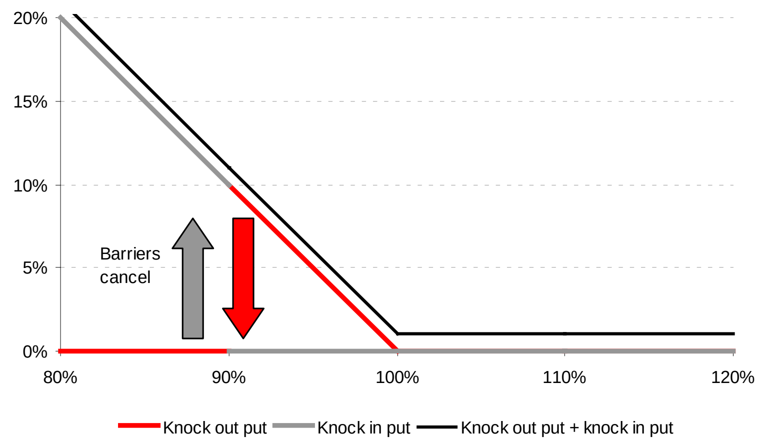

Source: Santander Investment Bolsa.

KNOCK-OUTS COST C15%-25% OF PUT SPREAD COST

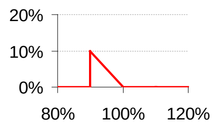

The payout of a knock-out put is equal to a 'shark fin' (see figure on the left) until the barrier is reached. A 'shark fin' is equal to a short digital position (at the barrier) plus a put spread (long put at strike of knock-out put, short put at barrier of knock-out put). The price of a knock-out put can therefore be considered to be the cost of a put spread, less a digital and less the value of the knock-out. As pricing digitals and barriers is not trivial, comparing the price of a knock-out put as a percentage of the appropriate put spread can be a quick way to evaluate value (the knock-out will have a lower value as it offers less payout to the downside). For reasonable barriers between 10% and 30% below the strike, the price of the knock-out option should be between c15% and c25% of the cost of the put spread.

CONTINUOUS BARRIERS ARE CHEAPER THAN DISCRETE

Continuous barriers are more popular as they are cheaper

There are two types of barriers, continuous and discrete. A continuous barrier is triggered if the price hits the barrier intraday, whereas a discrete barrier is only triggered if the closing price passes through the barrier. Discrete knock-out barriers are more expensive than continuous barriers, while the reverse holds for knock-in barriers (especially during periods of high volatility). There are also additional hedging costs to discrete barriers, as it is possible for spot to move through the barrier intraday without the discrete barrier being triggered (ie, if the close is the correct side of the discrete barrier). As these costs are passed on to the investor, discrete barriers are far less popular than continuous barriers for single stocks (c10%-20% of the market), although they do make up almost half the market for indices.

Jumps in stock prices between close and open is a problem for all barriers

While the hedge for a continuous barrier should, in theory, be able to be executed at a level close to the barrier, this is not the case should the underlying jump between close and open. In this case, the hedging of a continuous barrier suffers a similar problem to the hedging of a discrete barrier (delta hedge executed at a significantly different level to the barrier).

DOUBLE BARRIERS ARE POSSIBLE, BUT RARE

Double barrier options have both an up barrier and a down barrier. As only one of the barriers is significant for pricing, they are not common (as their pricing is similar to an ordinary single-barrier option). They make up less than 5% of the light exotic market.

REBATES CAN COMPENSATE INVESTORS IF BARRIER TRIGGERED

The main disadvantage of knock-out barrier options is that the investor receives nothing for purchasing the option if they are correct about the direction of the underlying (option is ITM) but incorrect about the magnitude (underlying passes through barrier). In order to provide compensation, some barrier options give the long investors a rebate if the barrier is triggered: for example, an ATM call with 120% knock-out that gives a 5% rebate if the barrier is touched. Rebates comprise approximately 20% of the index barrier market but are very rare for single-stock barrier options.

WORST-OF/BEST-OF OPTIONS

Worst-of (or best-of) options give payouts based on the worst (or best) performing asset. They are the second most popular light exotic due to structured product flow. Correlation is a key factor in pricing these options, and investor flow typically buys correlation (making uncorrelated assets with low correlation the most popular underlyings). The underlyings can be chosen from different asset classes (due to low correlation), and the number of underlyings is typically between three and 20. They are always European, and normally ATM options.

MATURITY IS NORMALLY ONE YEAR AND CAN BE CALLS OR PUTS

Worst-of/best-of options can be any maturity. Although the most popular is one-year maturity, up to three years can trade. As an option can be a call or a put, and either 'worst-of' or 'best-of'; there are four types of option to choose from. However, the most commonly traded are worst-of options (call or put). The payouts of the four types are given below:

Worst-of call payout = Max (Min ($r_1, r_2, \ldots, r_N$) , 0) where r_i is the return of N assets

Worst-of put payout = Max (-Min ($r_1, r_2, \ldots, r_N$) , 0) where r_i is the return of N assets

Best-of call payout = Max (Max ($r_1, r_2, \ldots, r_N$) , 0) where r_i is the return of N assets

Best-of put payout = Max (-Max ($r_1, r_2, \ldots, r_N$) , 0) where r_i is the return of N assets

WORST-OF CALLS POPULAR TO BUY (AS CHEAPER THAN ANY CALL)

The payout of a worst-of call option will be equal to the lowest payout of individual call options on each of the underlyings. As it is therefore very cheap, they are popular to buy. If all the assets are 100% correlated, then the value of the worst-of call is equal to the value of calls on all the underlyings (hence, in the normal case of correlation less than 100%, a worst-of call will be cheaper than any call on the underlying). If we lower the correlation, the price of the worst-of call also decreases (eg, the price of a worst-of call on two assets with -100% correlation is zero, as one asset moves in the opposite direction to the other). A worst-of call option is therefore long correlation. As worst-of calls are cheap, investors like to buy them and, therefore, provide buying pressure to implied correlation.

Rumour of QE2 lifted demand for worst-of calls on cross assets

Cross asset worst-of/best-of options are popular when macro risks dominate

Before QE2 (second round of quantitative easing) was announced, there was significant buying flow for worst-of calls on cross assets. The assets chosen were all assets that were likely to be correlated should QE2 occur but that would normally not necessarily be correlated (giving attractive pricing). QE2 was expected to cause USD weakening (in favour of other G10 currencies like the JPY, CHF and EUR), in addition to lifting 'risk-on' assets, like equities and commodities. The buying of worst-of calls on these three assets would therefore be a cheap way to gain exposure to the expected movements of markets if quantitative easing was extended (which it was).

WORST-OF PUTS ARE EXPENSIVE AND USUALLY SOLD

Flow from worst-of/best-of lifts implied correlation

A worst-of put will have a greater value than any of the puts on the underlying assets and is therefore very expensive to own. However, as correlation increases towards 100%, the value of the worst-of put will decrease towards the value of the most valuable put on either of the underlyings. A worst-of put is therefore short correlation. As selling (expensive) worst-of puts is popular, this flow puts buying pressure on implied correlation (the same effect as the flow for worst-of calls).

BEST-OF CALLS AND BEST-OF PUTS ARE RELATIVELY RARE

While worst-of options are popular, there is relatively little demand for best-of options. There are some buyers of best-of puts (which again supports correlation); however, best-of calls are very rare. Figure 81 below summarises the popularity and direction of investor flows (normally from structured products) and the effect on implied correlation. A useful rule of thumb for worst-of/best-of options is that they are short correlation if the price of the option is expensive (worst-of put and best-of call) and the reverse if the price of the option is cheap. This is why the buying of cheap and selling of expensive worst-of/best-of options results in buying flow to correlation.

Figure 81. Best-of/Worst-of Options

Option	Correlation	Flow	Cost	Notes
Worst-of put	Short	Sellers	Expensive	Popular structure to sell as cost is greater than that of most expensive put
Worst-of call	Long	Buyers	Cheap	Popular way to buy upside as low cost is less than cheapest call on any of the assets
Best-of put	Long	Some buyers	Cheap	Some buyers as cost is lower than cheapest put
Best-of call	Short	Rare	Expensive	Benefits from correlation falling as markets rise

Source: Santander Investment Bolsa.

LIGHT EXOTIC OPTIONS FLOW LIFTS IMPLIED CORRELATION

As the flow from worst-of/best-of products tends to support the levels of implied correlation, implied correlation typically trades above fair value. While other light exotic flow might not support correlation (eg, outperformance options, which are described below), worst-of/best-of options are the most popular light exotic, whose pricing depends on correlation and are therefore the primary driver for this market. We would point out that the most popular light exotics – barrier options – have no impact on correlation markets. In addition, worst-of/best-of flow is concentrated in uncorrelated assets, whereas outperformance options are usually on correlated assets.

OUTPERFORMANCE OPTIONS

Outperformance options are an option on the difference between returns on two different underlyings. They are a popular method of implementing relative value trades, as their cost is usually cheaper than an option on either underlying. The key unknown parameter for pricing outperformance options is implied correlation, as outperformance options are short correlation. The primary investor base for outperformance options is hedge funds, which are usually buyers of outperformance options on two correlated assets (to cheapen the price). Outperformance options are European and can always be priced as a call. Unless they are struck with a hurdle, they are an ATM option.

OUTPERFORMANCE OPTIONS ARE USUALLY SHORT-DATED CALLS

Outperformance options give a payout based on the difference between the *returns* of two underlyings. While any maturity can be used, they tend to be for maturities up to a year (maturities less than three months are rare). The payout formula for an outperformance option is below – by convention always quoted as a call of ‘r_A over r_B’ (as a put of ‘r_A over r_B’ can be structured as a call on ‘r_B over r_A’). Outperformance options are always European (like all light exotics) and are traded OTC.

Payout = Max ($r_A - r_B$, 0) where r_A and r_B are the returns of assets A and B, respectively

OPTIONS USUALLY ATM, CAN HAVE HURDLE AND ALLOWABLE LOSS

While outperformance options are normally structured ATM, they can be cheapened by making it OTM through a hurdle or by allowing an allowable loss at maturity (which simply defers the initial premium to maturity). While outperformance options can be structured ITM by having a negative hurdle, as this makes the option more expensive, this is rare. The formula for outperformance option payout with these features is:

Payout = Max ($r_A - r_B$ – hurdle, – allowable loss)

OUTPERFORMANCE OPTIONS ARE SHORT CORRELATION

The pricing of outperformance options depends on both the volatility of the two underlyings and the correlation between them. As there tends to be a more liquid and visible market for implied volatility than correlation, it is the implied correlation that is the key factor in determining pricing. Outperformance options are short correlation, which can be intuitively seen as: the price of an outperformance option must decline to zero if one assumes correlation rises towards 100% (two identical returns give a zero payout for the outperformance option).

As flow is to the buy side, some hedge funds outperformance call overwrite

Outperformance options are ideal for relative value trades

Outperformance options are ideal for implementing relative value trades, as they benefit from the upside, but the downside is floored to the initial premium paid. The primary investor base for outperformance options are hedge funds. While flow is normally to the buy side, the overpricing of outperformance options due to this imbalance has led some hedge funds to call overwrite their relative value position with an outperformance option.

MARGRABE'S FORMULA CAN BE USED FOR PRICING

An outperformance option volatility σ_{A-B} can be priced using Margrabe's formula given the inputs of the volatilities σ_A and σ_B of assets A and B, respectively, and their correlation ρ. This formula is shown below.

$$\sigma_{A-B} = \sqrt{\sigma_A{}^2 + \sigma_B{}^2 - 2\rho\sigma_A\sigma_B}$$

TEND TO BE USED FOR CORRELATED ASSETS

The formula above confirms mathematically that outperformance options are short correlation (due to the negative sign of the final term with correlation ρ). From an investor perspective, it therefore makes sense to sell correlation at high levels; hence, outperformance options tend to be used for correlated assets (so cross-asset outperformance options are very rare). This is why outperformance options tend to be traded on indices with a 60%-90% correlation and on single stocks that are 30%-80% correlated. The pricing of an outperformance option offer tends to have an implied correlation 5% below realised for correlations of c80%, and 10% below realised for correlations of c50% (outperformance option offer is a bid for implied correlation).

Best pricing is with assets of similar volatility

The price of an outperformance is minimised if volatilities σ_A and σ_B of assets A and B are equal (assuming the average of the two volatilities is kept constant). Having two assets of equal volatility increases the value of the final term $2\rho\sigma_A\sigma_B$ (reducing the outperformance volatility σ_{A-B}).

LOWER FORWARD FLATTERS OUTPERFORMANCE PRICING

ATMf strikes should be used when comparing relative costs

Assuming that the two assets have a similar interest rate and dividends, the forwards of the two assets approximately cancel each other out, and an ATM outperformance option is also ATMf (ATM forward or At The Money Forward). When comparing relative costs of outperformance options with call options on the individual underlyings, ATMf strikes must be used. If ATM strikes are used for the individual underlyings, the strikes will usually be lower than ATMf strikes and the call option will appear to be relatively more expensive compared to the ATMf (= ATM) outperformance option.

Pricing of ATM outperformance options is usually less than ATMf on either underlying

If two assets have the same volatility ($\sigma_A = \sigma_B$) and are 50% correlated (ρ = 50%), then the input for outperformance option pricing σ_{A-B} is equal to the volatilities of the two underlyings ($\sigma_{A-B} = \sigma_A = \sigma_B$). Hence, ATMf (ATM forward) options on either underlying will be the same as an ATMf (=ATM) outperformance option. As outperformance options tend to be used on assets with higher than 50% correlation and whose volatilities are similar, outperformance options are usually cheaper than similar options on either underlying.

LOOK-BACK OPTIONS

There are two types of look-back options, strike look-back and payout look-back, and both are usually multi-year options. Strike reset (or look-back) options have their strike set to the highest, or lowest, value within an initial look-back period (of up to three months). These options are normally structured so the strike moves against the investor in order to cheapen the cost. Payout look-back options conversely tend to be more attractive and expensive than vanilla options, as the value for the underlying used is the best historical value. As with all light exotics, these options are European and OTC.

STRIKE OF RESET OPTIONS MOVES AGAINST INVESTOR

Strike look-back, and payout look-back are the two types of look-back options

There are two main strike reset options, and both have an initial look-back period of typically one to three months, where the strike is set to be the highest (for a call) or lowest (for a put) traded value. While the look-back optionality moves against the investor, as the expiry of these options is multi-year (typically three), there is sufficient time for spot to move back in the investor's favour, and the strike reset cheapens the option premium. While having a strike reset that moves the strike to be the most optimal for the investor is possible, the high price means they are unpopular and rarely trade. While the cheaper form of strike reset options does attract some flow due to structured products, they are not particularly popular.

Strike reset options perform best when there is an initial period of range trading

There are three possible outcomes to purchasing a strike reset option. Strike reset options can be considered a cheaper alternative to buying an ATM option at the end of the strike reset period, as the strike is roughly identical for two of the three possible outcomes (but at a lower price).

- **Spot moves in direction of option payout.** If spot moves in a direction that would make the option ITM, the strike is reset to be equal to spot as it moves in a favourable direction, and the investor is left with a roughly ATM option.
- **Range-trading markets.** Should markets range trade, the investor will similarly receive a virtually ATM option at the end of the strike reset period.
- **Spot moves in opposite direction to option payout.** If spot initially moves in the opposite direction to the option payout (down for calls, up for puts), then the option strike is identical to an option that was initially ATM (as the key value of the underlying for the strike reset is the initial value) and, hence, OTM at the end of the strike rest period. The downside of this outcome is why strike reset options can be purchased for a lower cost than an ATM option.

Strike reset options are therefore most suitable for investors who believe there will be an initial period of range trading, before the underlying moves in a favourable direction.

PAYOUT LOOK-BACK OPTIONS

Having a look-back option that selects the best value of the underlying (highest for calls, lowest for puts) increases the payout of an option – and cost. These options typically have a five-year maturity and typically use end-of-month or end-of-year values for the selection of the optimal payout.

CONTINGENT PREMIUM OPTIONS

Contingent premium options are initially zero-premium and only require a premium to be paid if the option becomes ATM on the close. The contingent premium to be paid is, however, larger than the initial premium would be, compensating for the fact that it might never have to be paid. Puts are the most popular, giving protection with zero initial premium. These typically one-year put options are OTM (or the contingent premium would almost certainly have to be paid immediately) and European.

CONTINGENT PREMIUM OPTIONS ALLOW ZERO UPFRONT COST

Contingent premium allows protection to be bought at zero cost

While contingent premium calls are possible, the most popular form is for a contingent premium put to allow protection to be bought with no initial cost. The cost of the premium to be paid is roughly equal to the initial premium of the vanilla option, divided by the probability of spot trading through the strike at some point during the life of the option (eg, an 80% put whose contingent premium has to be paid if the underlying goes below 80%). Using contingent premium options for protection has the benefit that no cost is suffered if the protection is not needed, but if spot dips below the strike/barrier, then the large premium has to be paid (which is likely to be more than the put payout unless there was a large decline). These can be thought of as a form of 'crash put'.

Having a conditional premium on a level other than strike is possible, but rare

The usual structure for contingent premium options is to have the level at which the premium is paid equal to the strike. The logic is that although investors have to pay a large premium, they do have the benefit of holding an option that is slightly ITM. Having the conditional premium at a level other than strike is possible, but rare (eg, an 80% put whose contingent premium has to be paid if the underlying reaches 110%).

COMPOSITE AND QUANTO OPTIONS

There are two types of option involving different currencies. The simplest is a composite option, where the strike (or payoff) currency is in a different currency to the underlying. A slightly more complicated option is a quanto option, which is similar to a composite option, but the exchange rate of the conversion is fixed.

COMPOSITE OPTIONS USE DIFFERENT VOLATILITY INPUT

Composite options are usually more expensive than vanilla options

A composite option is a cash or physical option on a security whose currency is different from the strike or payoff currency (eg, Euro strike option on Apple). If an underlying is in a foreign currency, then its price in the payout (or strike) currency will usually be more volatile (and hence more expensive) due to the additional volatility associated with currency fluctuations. Only for significantly negative correlations will a composite option be less expensive than the vanilla option (if there is zero correlation the effect of FX still lifts valuations). The value of a composite option can be calculated using Black-Scholes as usual, by substituting the volatility of the asset with the volatility of the asset in payout currency terms. The payout (or strike) currency risk-free rate should be used rather than the (foreign) security currency risk-free rate. The dividend yield assumption is unchanged (as it has no currency) between a composite option and a vanilla option.

$$\sigma_{Payout} = \sqrt{\sigma_{Security}{}^2 + \sigma_{FX}{}^2 + 2\rho\sigma_{Security}\sigma_{FX}}$$

where

σ_{Payout} = volatility of asset in payout (strike) currency

$\sigma_{Security}$ = volatility of asset in (foreign) security currency

σ_{FX} = volatility of FX rate (between payout currency and security currency)

ρ = correlation of FX rate (security currency in payoff currency terms) and security price

Composite options are long correlation (if FX is foreign currency in domestic terms)

The formula to calculate the volatility of the underlying is given above. As the payoff increases with a positive correlation between FX and the underlying, a composite option is long correlation (the positive payout will be higher due to FX, while FX moving against the investor is irrelevant when the payout is zero). Note that care has to be taken when considering the definition of the FX rate; it should be the (foreign) security currency given in (domestic) payoff currency terms.

For example, if we are pricing a euro option on a dollar-based security and assume an extreme case of ρ = 100%, the volatility of the USD underlying in EUR will be the sum of the volatility of the underlying and the volatility of USD.

QUANTO OPTIONS USE DIFFERENT DIVIDEND INPUT

Quanto options cost a similar amount to vanilla options if FX correlation is small

Quanto options are similar to a composite option, except the payout is always cash settled and a fixed FX rate is used to determine the payout. Quanto options can be modelled using Black-Scholes. As the FX rate for the payout is fixed, quanto options are modelled using the normal volatility of the underlying (as FX volatility has no effect). The payout is simply the fixed FX rate multiplied by the price of a vanilla option with the same volatility, but a different carry. The carry (risk-free rate - dividend) to be used is shown below (the risk-free rate for quanto options is assumed to be the risk-free rate of the security currency, ie, it is not the same as for composite options).

$$c_{Quanto} = rfr_{Security} - d - \rho\sigma_{Security}\sigma_{FX}$$

⇨ $d_{Quanto} = d + \rho\sigma_{Security}\sigma_{FX}$ as $d_{quanto} = rfr_{Security} - c_{Quanto}$

where

c_{Quanto} = carry for quanto pricing

d_{Quanto} = dividend for quanto pricing

d = dividend yield

$rfr_{Security}$ = risk free rate of security currency

rfr_{Payout} = risk free rate of payout currency

Quanto options are either long or short correlation depending on the sign of the delta

The correlation between the FX and the security has an effect on quanto pricing, the direction (and magnitude) of which depends on the delta of the option. This is because the dividend risk of an option is equal to its delta, and the dividend used in quanto pricing increases as correlation increases.

Quanto option calls are short correlation (if FX is foreign currency in domestic terms)

As a call option is short dividends (call is an option on the price of underlying, not the total return of the underlying), a quanto call option is short correlation. A quanto put option is therefore slightly long correlation. In both cases, we assume the FX rate is the foreign security currency measured in domestic payout terms.

Intuitively, we can see a quanto call option is short correlation by assuming the dividend yield and both currency risk-free rates are all zero and comparing its value to a vanilla call option priced in the (foreign) security currency. If correlation is high, the vanilla call option is worth more than the quanto call option (as FX moves in favour of the investor if the price of the security rises). The reverse is also true (negative correlation causes a vanilla call option to be worth less than a quanto call option). As the price of a vanilla (single currency) call does not change due to the correlation of the underlying with the FX rate, this shows a quanto call option is short correlation.

This page left blank intentionally.

ADVANCED VOLATILITY TRADING

RELATIVE VALUE TRADING

Relative value is the name given to a variety of trades that attempt to profit from the mean reversion of two related assets that have diverged. The relationship between the two securities chosen can be fundamental (different share types of same company or significant cross-holding) or statistical (two stocks in same sector). Relative value can be carried out via cash (or delta-1), options or outperformance options.

TRADES ARE USUALLY CHOSEN ON CORRELATED ASSETS

Relative value trades attempt to profit from mean reversion

The payout of a relative value trade on two uncorrelated securities is completely random, and the investor on average gains no benefit. However, if two securities have a strong fundamental or statistical reason to be correlated, they can be thought of as trading in a similar direction with a random noise component. Assuming the correlation between the securities is sufficiently strong, the noise component should mean revert. Relative value trades attempt to profit from this mean reversion. There are five main types of relative value trades.

- **Dual listing.** If a share trades on different exchanges (eg, an ADR), the two prices should be equal. This is not always the case due to execution risk (different trading times) and perhaps due to indexation flow. Non-fungible shares or those with shorting restrictions are most likely to show the largest divergence in price. Of all relative value trades, dual-listing ones are likely to show the strongest correlation.

- **Share class.** If there is more than one type of share, perhaps with voting or ownership restrictions, then the price of these shares can diverge from one another. For example, preference shares typically have a higher dividend to compensate for lack of voting rights, but suffer from less liquidity and (normally) exclusion from equity indices. During special situations, for example, during the Porsche/VW saga, the difference in price between the two shares can diverge dramatically.

- **Cross-holding.** If one company (potentially a holding company) owns a significant amount of another company, the prices of the two companies will be linked. Sometimes putting on a cross-holding trade is difficult in practice due to the high borrow cost of the smaller company. This trade is also known as a stub trade when the investor wants pure exposure to the larger company, and hedges out the unwanted exposure to the equity holdings of the larger company. Potentially, these trades can occur when a larger company spins off a subsidiary but keeps a substantial stake post spin-off.

- **Event-driven.** In the event of a takeover that is estimated to have a significant chance of succeeding, the share prices of the acquiring and target company should be correlated. The target will usually trade at a discount to the bid price, to account for the probability the deals falls through (although if the offer is expected to be improved, or beaten by another bidder, the target could trade above the offer price).

- **Long-short.** If a long and short position is initiated in two securities that do not have one of the above four reasons to be correlated, it is a long-short trade. The correlation between the two securities of a long-short trade is likely to be lower than for other relative values trades. Because of this, often two stocks within a sector are chosen, as they should have a very high correlation and the noise component is likely to be bounded (assuming market share and profitability is unlikely to change substantially over the period of the relative value trade).

Long-short can focus returns on stock picking ability (which is c10% of equity return)

Only c20% of equity returns are due to stock picking (c10% is sector selection and c70% is the broad equity market)

General market performance is typically responsible for c70% of equity returns, while c10% is due to sector selection and the remaining c20% due to stock picking. If an investor wishes to focus returns on the proportion due to sector or stock picking, they can enter into a long position in that security and a short position in the appropriate market index (or vice versa). This will focus returns on the c30% due to sector and stock selection. Typically, relatively large stocks are selected, as their systematic risk (which should cancel out in a relative value trade) is usually large compared to specific risk. Alternatively, if a single stock in the same sector (or sector index) is used instead of the market index, then returns should be focused on the c20% due to stock picking within a sector.

SIZE OF POSITIONS SHOULD BE WEIGHTED BY BETA

If the size of the long-short legs are chosen to have equal notional (share price × number of shares × FX), then the trade will break even if both stock prices go to zero. However, the legs of the trade are normally weighted by beta to ensure the position is market neutral for more modest moves in the equity market. The volatility (historical or implied) of the stock divided by the average volatility of the market can be used as an estimate of the beta.

DELTA-1, OPTIONS AND OUTPERFORMANCE OPTIONS

Relative value trades can be implemented via cash/delta-1, vanilla options or outperformance options. They have very different trade-offs between liquidity and risk. No one method is superior to others, and the choice of which instrument to use depends on the types of liquidity and risk the investor is comfortable with.

Figure 82. Different Methods of Relative Value Trading

Asset Class	Position	Benefits	Disadvantages
Cash/delta-1	Long A, short B using stock/CFD, future, forwards, total return swap or ETF	High liquidity (volatility products might not be available)	Unlimited risk
Options	Long call on A, short call on B (or put/call spread/put spread)	Limited downside on long leg and convex payoff	Unlimited risk on short side (unless call spreads/put spreads)
Outperformance option	Long outperformance option on A vs B	Limited downside and convex payoff	Poor liquidity/wide bid-offer spreads

Source: Santander Investment Bolsa estimates.

(1) CASH/DELTA-1: BEST LIQUIDITY, BUT UNLIMITED RISK

The deepest and most liquid market for relative value trades is the cash (or delta-1) market. While there are limited restrictions in the size or stocks available, the trade can suffer potentially unlimited downside. While there are many similarities between cash or delta-1 instruments, there are also important differences.

Benefits of more beneficial taxation can be shared

For many delta-1 products, the presence of investors with more beneficial taxation can result in more competitive pricing. Products that have to be based in one location, such as ETFs, suffer from being unable to benefit from the different taxation of other investors.

Figure 83. Delta-1 Product Summary

	Listed / OTC	Counter party risk	Leverage	Maintenance	Benefit from borrow	Dividend risk	Cost	Notes	Overall
ETFs	Listed	Low	No (cash)	Low	No	No	Medium	Dividend taxation	
Futures	Listed	Low	Yes	Medium	Yes	Yes	Medium	Tracking error	
Forwards	OTC	Medium	Yes	Low	Yes	Yes	Low	Rare	
Total return swaps	OTC	Medium	Yes	Low	Yes	No	Low	Best OTC	
CFDs	OTC	Medium	Yes	Medium	No	No	Low	Similar to TRS	
Cash	Listed	Low	No (cash)	High	No	No	Low	Keep voting rights	
Certificates / p-notes	Listed	High (issuer)	No (cash)	Low	No	No	High	Flexible underlying	

Source: Santander Investment Bolsa.

CFDs have many advantages over stock

A relative value trade in the cash (stock) market can be initiated by a long stock position combined with a short stock position. The short stock position needs a functioning stock borrow market, as the stock needs to be borrowed before it can be sold short. Using stock can tie up a lot of capital, as the long position needs to be funded, as does the short position. Normally, a prime broker can help fund the position; however, for simplicity, CFDs are often used instead of stock.

As CFDs remove the overhead of corporate actions such as dividends, they are very popular with hedge funds that wish to quickly initiate long/short positions. As CFDs avoid paying stamp duty in certain countries, there can be yield enhancement benefits from using CFDs. While in theory a CFD has counterparty risk, a CFD is often created to be a stock equivalent with daily resets (exchange of cash flows) limiting this disadvantage. The main disadvantage of using CFDs is the loss of voting rights; however, relative value investors are not usually interested in voting. While stock and CFDs can be used to trade indices, this is rare as it usually requires more maintenance than other delta-1 products.

Total return swaps and forwards are best for indices

Total return swap is the best instrument for trading an index

The index equivalent of a CFD is a total return swap. Potentially, a portfolio swap could be used instead of a total return swap to manage the long and short legs in one trade. A forward is essentially a price return equivalent of a total return swap and is normally only used instead of a TRS (total return swap) for internal reasons (for example, if IT systems can only handle forwards). Both futures and forwards benefit from the more optimal taxation treatment of other investors, allowing yield enhancement.

Futures and ETFs provide a listed method of trading indices with no counterparty risk

Future is the best listed (no counterparty risk) instrument for trading an index

A future is a listed equivalent of a forward (hence, it suffers from exchange fees) and benefits from the same yield enhancement factors. Futures often require slightly larger margins to take into consideration there is no counterparty risk. For investors requiring a listed instrument with no counterparty risk, futures are the best instrument for trading indices. There is, however, a significant maintenance cost to futures, as liquidity is concentrated on the front months and therefore requires rolling. While futures and forwards can be used to trade single stocks, they are usually used for indices (although in Europe futures can be crossed, allowing them to be used instead of stock). An ETF can be thought of as the index equivalent of a stock, being a fully funded listed instrument with no potential for yield enhancement. As there are more ETFs than futures, they can be used for a wider variety of underlyings.

Figure 84. Delta-1 Product Decision Tree

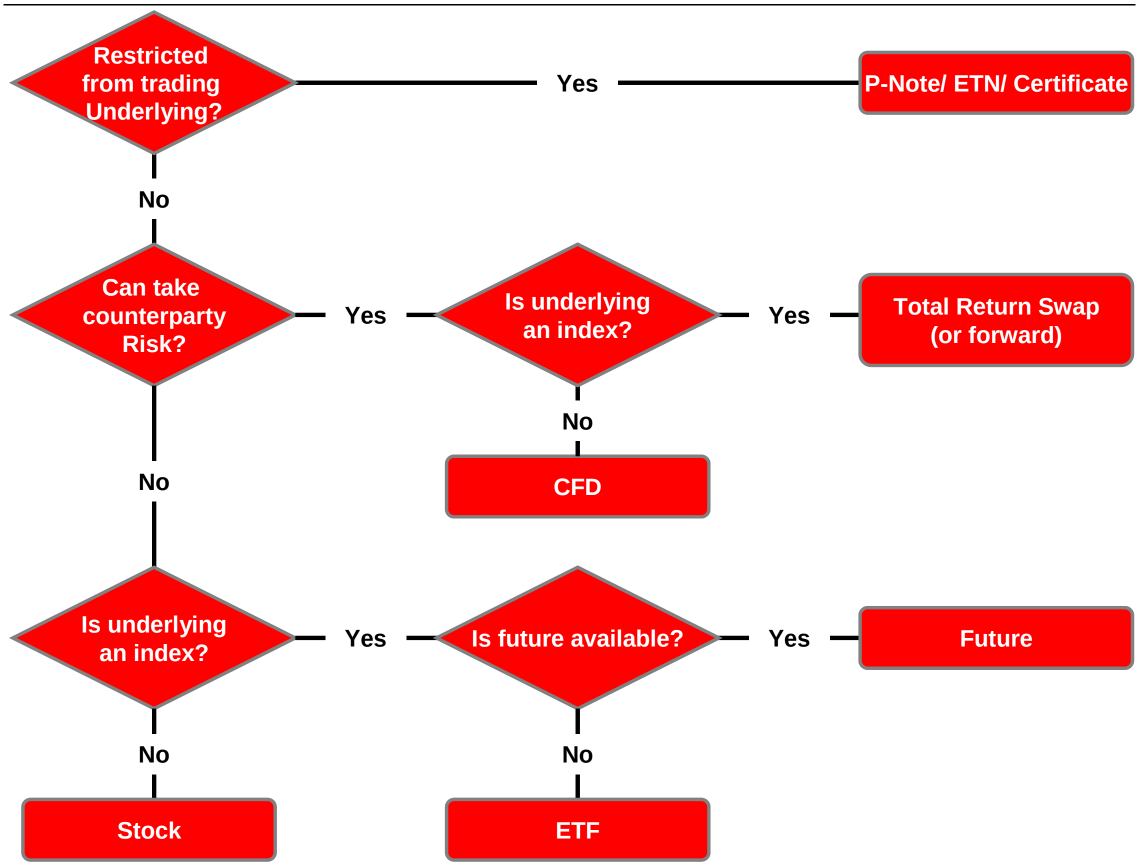

Source: Santander Investment Bolsa.

Certificates, ETNs and p-notes are effectively the same

Certificates/ETNs and participation notes (p-notes) are traded on an exchange and can give exposure to strategies, markets and currencies that an investor might normally be unable to invest in. For example, if an investor is prohibited from investing in volatility then a certificate or ETN that wraps a volatility strategy (call overwriting, selling one-month variance swaps, long VIX/vStoxx futures, etc) can be bought instead. Access to Chinese and Indian markets is not trivial, but can be traded via p-notes (as can trading in markets with restricted currencies, as the product can be redenominated in USD or another currency). As 100% upfront payment is required, certificate/ETNs and p-notes can be considered a fully funded equity swap (total or price return) with a listed price. Despite being listed, the investor is 100% exposed to the credit risk of the counterparty.

(2) OPTIONS: CONVEX PAYOFF AND CAN LIMIT DOWNSIDE ON LONG LEG

Long leg of relative value trade is often rotated into options when long leg becomes profitable

Options can be used in place of stock or delta-1 for either the long or short leg, or potentially both. Options offer convexity, allowing a position to profit from the expected move while protecting against the potentially unlimited downside. Often a relative value trade will be put on in the cash/delta-1 market, and the long leg rotated into a call once the long leg is profitable (in order to protect profits). While volatility is a factor in determining the attractiveness of using options, the need for safety or convexity is normally the primary driver for using options (as relative value traders do not delta hedge, the change in implied volatility is less of a factor in profitability than the delta/change in equity market). Investors who are concerned about the cost of options can cheapen the trade by using call spreads or put spreads in place of vanilla calls or puts.

Weighting options by volatility is similar to weighting by beta and roughly zero cost

The most appropriate weighting for two relative value legs is beta weighting the size of the delta hedge of the option (ie, same beta × number of options × delta × FX), rather than having identical notional (share price × number of options × FX). Beta weighting ensures the position is market neutral. Volatility weighting can be used as a substitute for beta weighting, as volatility divided by average volatility of the market is a reasonable estimate for beta. Volatility weighting ATM (or ATMf) options is roughly zero cost, as the premium of ATM options is approximately linear in volatility.

Choosing strike and maturity of option is not trivial

One disadvantage of using options in place of equity is the need to choose a maturity, although some investors see this as an advantage as it forces a view to be taken on the duration or exit point of the trade at inception. If the position has to be closed or rolled before expiry, there are potentially mark-to-market risks. Similarly, the strike of the option needs to be chosen, which can be ATM (at the money), ATMf (ATM forward), same percentage of spot/forward or same delta. Choosing the same delta of an OTM option means trading a strike further away from spot/forward for the more volatile asset (as delta increases as volatility increases). We note that trading the same delta option is not the same as volatility weighting the options traded as delta is not linear in volatility. Delta also does not take into account the beta of the underlyings.

(3) OUTPERFORMANCE OPTIONS: LIMITED DOWNSIDE BUT LOW LIQUIDITY

Outperformance options are ideally suited to relative value trades, as the maximum loss is the premium paid and the upside is potentially unlimited. However, outperformance options suffer from being relatively illiquid. While pricing is normally cheaper than vanilla options (for normal levels of correlation), it might not be particularly appealing depending on the correlation between the two assets. As there are usually more buyers than sellers of outperformance options, some hedge funds use outperformance options to overwrite their relative value trades.

RELATIVE VALUE VOLATILITY TRADING

Volatility investors can trade volatility pairs in the same way as trading equity pairs. For indices, this can be done via options, variance swaps or futures on a volatility index (such as the VIX or vStoxx). For indices that are popular volatility trading pairs, if they have significantly different skews this can impact the volatility market. Single-stock relative value volatility trading is possible, but less attractive due to the wider bid-offer spreads.

THERE ARE TWO WAYS TO PROFIT FROM VOLATILITY PAIR TRADING

When a pair trade between two equities is attempted, the main driver of profits is from a mean reversion of the equity prices. With volatility relative value trading, there are two ways of profiting:

- **Mean reversion.** In the same way an equity pair trade profits from a mean reversion of stock prices, a volatility pair trade can profit from a mean reversion of implied volatility. For short-term trades, mean reversion is the primary driver for profits (or losses). For relative value trades using forward starting products (eg, futures on volatility indices), this is the only driver of returns as forward starting products have no carry. The method for finding suitable volatility pair trades that rely on a short-term mean reversion are similar to that for a vanilla pair trade on equities.
- **Carry.** For an equity pair trade, the carry of the position is not as significant as, typically, the dividend yields of equities do not differ much from one another and are relatively small compared to the movement in spot. However, the carry of a volatility trade (difference between realised volatility and implied volatility) can be significant. As the duration of a trade increases, the carry increases in importance. Hence, for longer term volatility pair trades it is important to look at the difference between realised and implied volatility.

IMPLIED VOLATILITY SPREAD BETWEEN PAIRS IS KEPT STABLE

Relative value volatility trades can impact the volatility market

While the skew of different indices is dependent on correlation, traders tend to keep the absolute difference in implied volatility stable due to mean reversion. This is why if equity markets move down, the implied volatility of the S&P500 or FTSE (as they are large diversified indices that hence have high skew) tends to come under pressure, while the implied volatility of country indices with fewer members, such as the DAX, are likely to be supported. The SX5E tends to lie somewhere in between, as it has fewer members than the S&P500 or FTSE but is more diverse than other European country indices. Should markets rise, the reverse tends to occur (high skew indices implieds are lifted, low skew implieds are weighed on).

Difference between implieds is key, not the absolute level of each implied

We note that for returns due to mean reversion, it is not the absolute level of volatility that is key but the difference. For example, let us assume stock A implieds trade between 20% and 25% while stock B implieds trade between 30% and 35%. If stock A is at 25% implied (top of range) while stock B implied is at 30% implied (bottom of range), a short A volatility long B volatility position should be initiated. This is despite the 25% implied of A being less than the 30% implied of B.

VOLATILITY PAIR TRADING GREEKS ARE SIMILAR TO DISPERSION

Dispersion trading could be considered to be a pair trade where one leg is a basket of single stock volatility

In dispersion trading, a (normally short) index volatility position is traded against a basket of (normally long) single-stock volatility positions. This position of index volatility vs basket could be considered to be a pair trade where one leg is the index and the other leg is the basket. A pair trade can be carried out via straddles / strangles or variance swaps, just like dispersion. We shall assume that the pair trade is being carried out by delta hedging options, for trading via variance swaps simply replaces notional in the table below with the vega of the variance swap. The weighting of the legs in order to be vega / theta or gamma flat is similar to dispersion trading, as can be seen below.

Figure 85. Greeks of Option Pair Trades with Different Weightings (shorting low vol, long high vol)

Greeks	Theta-Weighted	Vega-Weighted	Dollar Gamma-Weighted
Theta	0	Pay (or negative/short)	Pay a lot (very negative/short)
Vega	Short	0	Long
Gamma	Very short	Short	0
Ratio high vol to low vol notional	σ low vol / σ high vol	1	σ high vol / σ low vol
Notional of high vol stock	**Less than low vol**	**Equal to low vol**	**More than low vol**

Source: Santander Investment Bolsa.

Sign of theta, vega and gamma depends on which way round the pair trade is initiated

The sign of theta, vega and gamma are based on a trade of shorting the lower volatility security and going long the higher volatility security (on an absolute basis) in order for easy comparison to dispersion trading (where, typically, the lower absolute volatility of the index is shorted against a long of the higher absolute volatility of the single stocks). For the reverse trade (short the higher absolute volatility security and long the lower absolute volatility security), the signs of the greeks need to be reversed.

PAIR TRADES CAN BE THETA OR VEGA WEIGHTED

Theta and vega weighted are the most common methods of weighting pair trades. Dollar gamma weighted is rarely used and is included for completeness purposes only. Theta-weighted trades assume proportional volatility changes (eg, if stock A has 20% implied and stock B has 25% implied, if stock A rises from 20% to 30% implied that is a 50% increase and stock B rises 50% to 37.5% implied). Vega-weighted trades assume absolute volatility changes (eg, if stock A has 20% implied and stock B has 25% implied, if stock A rises from 20% to 30% that is a 10 volatility point increase and stock B rises 10 volatility points to 35% implied).

Pair trade between two securities of same type should be theta weighted

If a pair trade between two securities of the same type (ie, two indices, or two single stocks) is attempted, theta weighting is the most appropriate. This is because the difference between a low volatility security and a high volatility security (of the same type) usually increases as volatility increases (ie, a proportional move). If a pair trade between an index and a single stock is attempted, vega weighting is the best as the implied volatility of an index is dependent not only on single-stock implied volatility but also on implied correlation. As volatility and correlation tend to move in parallel, this means the payout of a vega-weighted pair trade is less dependent on the overall level of volatility (hence the volatility mispricing becomes a more significant driver of the P&L of the trade)[25].

[25] There is evidence to suggest that vega-weighted index vs single-stock pair trades on average associate 2%-5% too much weight to the single-stock leg compared to the index leg. However, as this is so small compared to stock specific factors, it should be ignored.

TRADING EARNINGS ANNOUNCEMENTS/JUMPS

From the implied volatilities of near dated options, it is possible to calculate the implied jump on key dates. Trading these options in order to take a view on the likelihood of unanticipated (low or high) volatility on reporting dates is a very common strategy. We examine the different methods of calculating the implied jump, and show how the jump calculation should normalise for index term structure.

TOTAL VOLATILITY = DIFFUSIVE VOLATILITY + JUMP VOLATILITY

Expiry to trade jump is the expiry just after the event (can hedge by shorting expiry before event if available)

While stock prices under Black-Scholes are modelled as having a GBM (Geometric Brownian Motion) with constant volatility, in reality there are certain dates where there is likely to be more volatility than average. These key dates are usually reporting dates, but could also coincide with conference dates or investor days (in fact, any day where material non-public information is released to the public). The implied volatility of an option whose expiry is after a key date can be considered to be the sum of the normal diffusive volatility (normal volatility for the stock in the absence of any significantly material information being released) and the volatility due to the anticipated jump on the key date. While options of any expiry after the key date could be used, we shall assume the expiry chosen is the expiry just after the key date (to ensure the greatest percentage of the options' time value is associated with the jump). This position can be hedged by shorting the expiry before the key date, if one exists.

ESTIMATING DIFFUSIVE VOLATILITY IS NOT TRIVIAL

In order to calculate the implied jump due to a key date, the diffusive (normal) volatility of the stock needs to be estimated. While the diffusive volatility could be estimated by looking at historical volatility, it is usual to look at implied volatility (as there are several measures of historical volatility, but only one implied volatility). If there is an option that expires just before the key date, then the implied volatility of this option can be used. If not, the forward volatility after the key date is used as the estimate for the normal volatility of the security.

Figure 86. Diffusive Assumption Using Implied Vol **Diffusive Assumption Using Forward Vol**

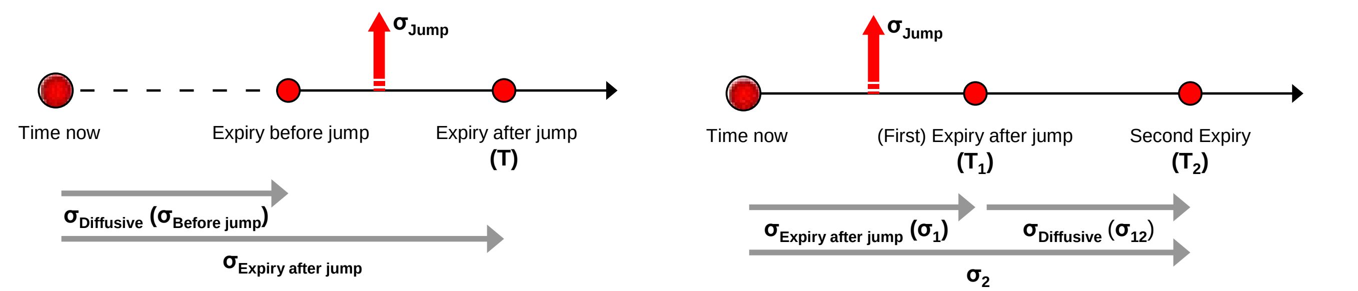

Source: Santander Investment Bolsa estimates.

Implied jumps normally calculated for near-dated events

Implied jumps are normally only calculated for near-dated events, as the effect of the jump tends to be too diluted for far dated expiries (and hence would be untradeable taking bid-offer spreads into account). Forward starting options could be used to trade jumps after the first expiry, but the wider bid-offer spread could be greater than potential profits.

Forward volatility can be calculated with implied of two options

The calculation for forward volatility is derived from the fact variance (time weighted) is additive. The formula is given below (σ_x is the implied volatility for options of maturity T_x).

$$\sigma_{12} = \sqrt{\frac{\sigma_2{}^2 T_2 - \sigma_1{}^2 T_1}{T_2 - T_1}} = \text{forward volatility } T_1 \text{ to } T_2$$

JUMP VOLATILITY CAN BE CALCULATED FROM DIFFUSIVE VOLATILITY

Volatility just after reporting tends to be ¾ of the volatility just before reporting

As variance is additive, the volatility due to the jump can be calculated from the total volatility and the diffusive volatility. We note this assumes the normal diffusive volatility is constant, whereas volatility just after a reporting date is, in fact, typically ¾ of the volatility just before a reporting date (as the previously uncertain earnings are now known).

$$\sigma_{\text{Expiry after jump}}{}^2 T = \sigma_{\text{Jump}}{}^2 + \sigma_{\text{Diffusive}}{}^2 (T-1)$$

$$\Rightarrow \quad \sigma_{\text{Jump}} = \sqrt{(\sigma_{\text{Expiry after jump}}{}^2 T - \sigma_{\text{Diffusive}}{}^2 (T-1)}$$

where

$\sigma_{\text{Expiry after jump}}$= implied volatility of option whose expiry is after the jump

T = time to the expiry after jump (= T_1)

$\sigma_{\text{Diffusive}}$ = diffusive volatility ($\sigma_{\text{Before jump}}$ if there is an expiry before the jump, if not it is σ_{12})

σ_{Jump}= implied volatility due to the jump

IMPLIED JUMP CALCULATED FROM JUMP VOLATILITY

From the above implied volatility due to jump (σ_{Jump}) it is possible to calculate the implied daily return on the day of the jump (which is a combination of the normal daily move and the effect of the jump).

$$\text{Expected daily return} = e^{\frac{\sigma_{\text{Jump}}{}^2}{2}} \left[(2 \times N(\sigma_{\text{Jump}}) - 1) \right]$$

A proof of the above formula is below.

$$\Delta S = S_1 - S_0 = S_0 \left(\frac{S_1}{S_0} - 1 \right) = S_0 \left(e^r - 1 \right) \text{ where } r = Ln\left(\frac{S_1}{S_0} \right)$$

$$\Rightarrow \quad \mathrm{E}\left(\frac{|\Delta S|}{S_0} \right) = \mathrm{E}\left(\left| e^r - 1 \right| \right)$$

$$\Rightarrow \quad \text{Expected daily return} = \mathrm{E}\left(\left| e^r - 1 \right| \right)$$

as r is normally distributed

$$\Rightarrow \text{Expected daily return} = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} \left(e^r - 1\right) e^{\frac{-r^2}{2\sigma^2}} dr$$

$$\Rightarrow \text{Expected daily return} = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{0}^{\infty} \left(e^r - 1\right) e^{\frac{-r^2}{2\sigma^2}} dr + \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{0} \left(e^r - 1\right) e^{\frac{-r^2}{2\sigma^2}} dr$$

$$\Rightarrow \text{Expected daily return} = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{0}^{\infty} \left(e^r - e^{-r}\right) e^{\frac{-r^2}{2\sigma^2}} dr$$

if we define x such that r = xσ

$$\Rightarrow \text{Expected daily return} = \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} \left(e^{x\sigma} - e^{-x\sigma}\right) e^{\frac{-x^2}{2}} dx$$

$$\Rightarrow \text{Expected daily return} = \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{x\sigma - \frac{x^2}{2}} - e^{-x\sigma - \frac{x^2}{2}} dx$$

$$\Rightarrow \text{Expected daily return} = \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{x\sigma - \frac{x^2}{2}} dx - \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{-x\sigma - \frac{x^2}{2}} dx$$

$$\Rightarrow \text{Expected daily return} = \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{\frac{-(x-\sigma)^2}{2} + \frac{\sigma^2}{2}} dx - \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{\frac{-(x+\sigma)^2}{2} + \frac{\sigma^2}{2}} dx$$

$$\Rightarrow \text{Expected daily return} = e^{\frac{\sigma^2}{2}} \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{\frac{-(x-\sigma)^2}{2}} dx - e^{\frac{\sigma^2}{2}} \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{\frac{-(x+\sigma)^2}{2}} dx$$

$$\Rightarrow \text{Expected daily return} = e^{\frac{\sigma^2}{2}} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{0} e^{\frac{-(x-\sigma)^2}{2}} dx - e^{\frac{\sigma^2}{2}} \frac{1}{\sqrt{2\pi}} \int_{0}^{\infty} e^{\frac{-(x+\sigma)^2}{2}} dx$$

$$\Rightarrow \text{Expected daily return} = e^{\frac{\sigma^2}{2}} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\sigma} e^{\frac{-x^2}{2}} dx - e^{\frac{\sigma^2}{2}} \frac{1}{\sqrt{2\pi}} \int_{\sigma}^{\infty} e^{\frac{-x^2}{2}} dx$$

$$\Rightarrow \text{Expected daily return} = e^{\frac{\sigma^2}{2}} N(\sigma) - e^{\frac{\sigma^2}{2}} \left[1 - N(\sigma)\right]$$

$$\Rightarrow \text{Expected daily return} = e^{\frac{\sigma^2}{2}} \left[2 \times N(\sigma) - 1\right]$$

EQUITY TERM STRUCTURE CAN BE ADJUSTED BY INDEX TERM STRUCTURE

The methodology for extracting jumps from the difference between the front-month implieds is simply a case of mathematics, assuming the volatility of a stock is equal to a 'normal' volatility on every day plus an 'abnormal' jump on reporting. In order to calculate the 'abnormal' jump, we need to estimate the 'normal' volatility, and this estimate usually requires a flat term structure to be assumed. If the index term structure is used to adjust the single-stock term structure, then a more accurate implied jump can be calculated [26](assuming the single-stock term structure would be identical to index term structure without the effect of a reporting date). For simplicity, the diagrams below assume reporting is between expiry 2 and 3, but the effect will be similar if earnings is between expiry 1 and 2.

Figure 87. Equity and Index Term Structure **Equity Term Structure Adjusted by Index Term Structure**

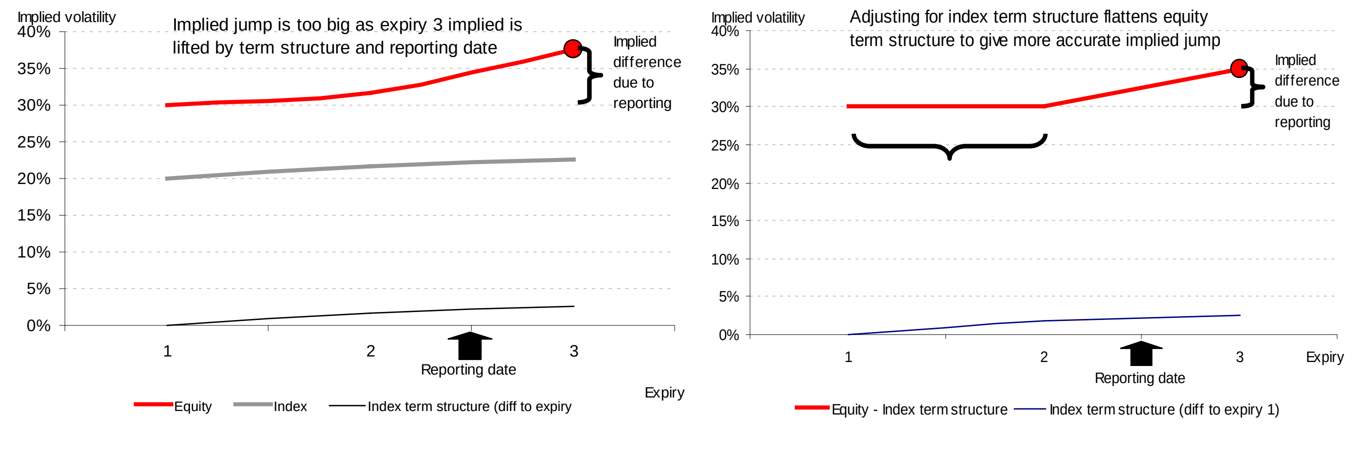

Source: Santander Investment Bolsa estimates.

[26] This assumes a flat implied correlation term structure, which is a reasonable assumption for the very near-dated expiries.

STRETCHING BLACK-SCHOLES ASSUMPTIONS

The Black-Scholes model assumes an investor knows the future volatility of a stock, in addition to being able to continuously delta hedge. In order to discover what the likely profit (or loss) will be in reality, we stretch these assumptions. If the future volatility is unknown, the amount of profit (or loss) will vary depending on the path, but buying cheap volatility will always show some profit. However, if the position is delta-hedged discretely, the purchase of cheap volatility may reveal a loss. The variance of discretely delta-hedged profits can be halved by hedging four times as frequently. We also show why traders should hedge with a delta calculated from expected – not implied – volatility, especially when long volatility.

BLACK-SCHOLES ASSUMES KNOWN VOL AND CONTINUOUS HEDGING

While there are a number of assumptions behind Black-Scholes, the two which are the least realistic are: (1) a continuous and known future realised volatility; and (2) an ability to delta hedge continuously. There are, therefore, four different scenarios to investigate. We assume that options are European (can only be exercised at maturity), although most single-stock options are American (can be exercised at any time).

- **Continuous delta hedging with known volatility.** In this scenario, the profit (or loss) from volatility trading is fixed. If the known volatility is constant, then the assumptions are identical to Black-Scholes. Interestingly, the results are the same if volatility is allowed not to be constant (while still being known).

- **Continuous delta hedging with unknown volatility.** If volatility is unknown, then typically traders hedge with the delta calculated using implied volatility. However, as implied volatility is not a perfect predictor of future realised volatility, this causes some variation in the profit (or loss) of the position. However, with these assumptions, if realised volatility is above the implied volatility price paid, it is impossible to suffer a loss.

- **Discrete delta hedging with known volatility.** As markets are not open 24/7, continuous delta hedging is arguably an unreasonable assumption. The path dependency of discrete delta hedging adds a certain amount of variation in profits (or losses), which can cause the purchase of cheap volatility (implied less than realised) to suffer a loss. The variance of the payout is inversely proportional to the frequency of the delta hedging. For example, the payout from hedging four times a day has a variance that is a quarter of the variance that results if the position is hedged only once a day. The standard deviation is therefore halved if the frequency of hedging is quadrupled (as standard deviation squared = variance).

- **Discrete delta hedging with unknown volatility.** The most realistic assumption is to hedge discretely with unknown volatility. In this case, the payout of volatility trading is equal to the sum of the variance due to hedging with unknown volatility plus the variance due to discretely delta hedging.

CONTINUOUS DELTA HEDGING WITH KNOWN VOLATILITY

With continuous and known volatility, the correct delta can be calculated

In a Black-Scholes world, the volatility of a stock is constant and known. While a trader is also able to continuously delta hedge, Figure 88 below will assume we hedge discretely but in an infinitesimally small amount of time. In each unit of time, the stock can either go up or down. As the position is initially delta-neutral (ie, delta is zero), the gamma (or convexity) of the position gives it a profit for both downward and upward movements. While this effect is always profitable, the position does lose time value (due to theta). If an option is priced using the actual fixed constant volatility of the stock, the two effects cancel each other and the position does not earn an abnormal profit or loss as the return is equal to the risk-free rate. There is a very strong relationship between gamma and theta (theta pays for gamma)[27].

Figure 88. Constant and Known Realised Volatility to Calculate Delta

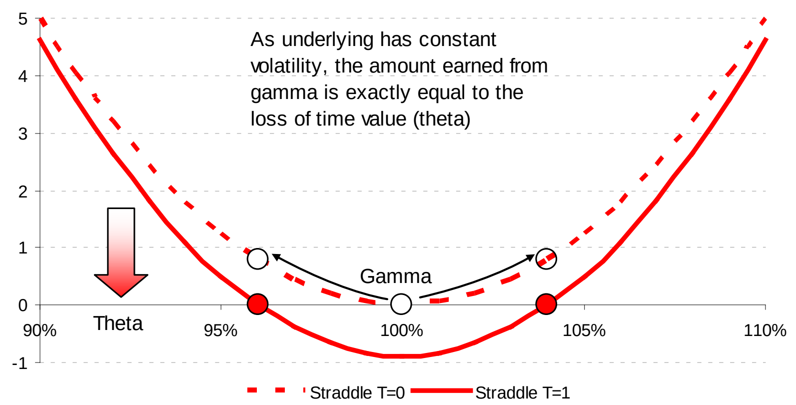

Source: Santander Investment Bolsa.

Profit from delta hedging is equal to the difference between price and theoretical price

The theoretical price of an option, using the known volatility, can be extracted by delta hedging. Should an option be bought at an implied volatility less than realised volatility, the difference between the theoretical price and the actual price will equal the profit of the trade. Figure 89 below shows the profit vs the difference in implied and realised volatility. As there is no path dependency, the profit (or loss) of the trade is fixed and cannot vary.

[27] They are not perfectly correlated, due to the interest paid on borrowing the shares (which varies with spot).

Figure 89. Profit (or Loss) from Continuously Delta-Hedging Known Volatility

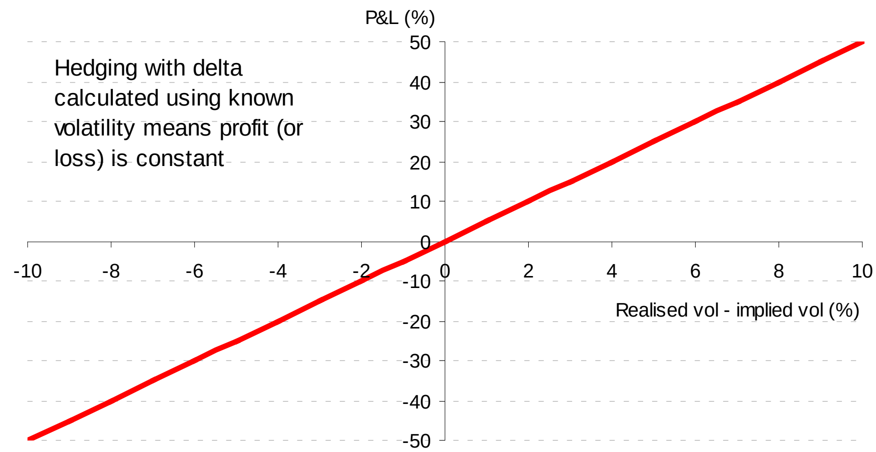

Source: Santander Investment Bolsa.

As theta and gamma are either both high or both low, profits are not path dependent

Delta hedging: an option that remains ATM and earns more gamma but pays more theta

If a position is continuously delta hedged with the correct delta (calculated from the known future volatility over the life of the option), then the payout is not path dependent. Figure 90 below shows two paths with equal volatility and the same start and end point. Even though one path is always ATM while the other has most volatility OTM, delta hedging gives the same profit for both. The cause of this relationship is the fact that, while the ATM option earns more due to delta hedging, the total theta cost is also higher (and exactly cancels the delta hedging profit).

Figure 90. Two Security Paths with Identical Volatility, Start and End Points

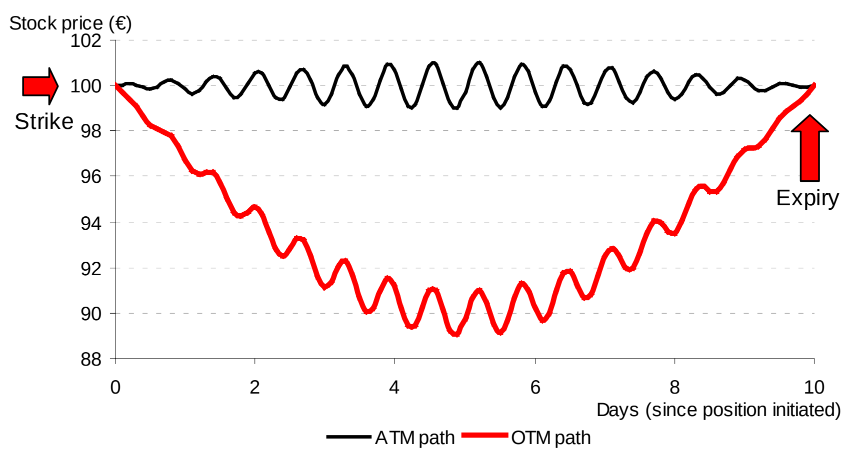

Source: Santander Investment Bolsa.

Profits are path independent, even if volatility is not constant (but still known)

While Black-Scholes assumes a constant known volatility, there are similar results for non-constant known volatility. This result is due to the fact that a European option payout depends only on the stock price at expiry. Therefore, the volatility over the life of the option is the only input to pricing. The timing of this volatility is irrelevant.

CONTINUOUS DELTA HEDGING WITH UNKNOWN VOLATILITY

If volatility is unknown, the correct delta cannot be calculated

As it is impossible to know in advance what the future volatility of a security will be, the implied volatility is often used to calculate deltas. Delta hedging using this estimate causes the position to have equity market risk and, hence, it becomes path dependent (although the average or expected profit remains unchanged). Figure 91 below shows that the profits from delta hedging are no longer independent of the direction in which the underlying moves. The fact that there is a difference between the correct delta (calculated using the remaining volatility to be realised over the life of the option) and the delta calculated using the implied volatility means returns are dependent on the direction of equity markets.

Figure 91. Profit from Cheap Options Is Not Constant if Volatility Is Not Known

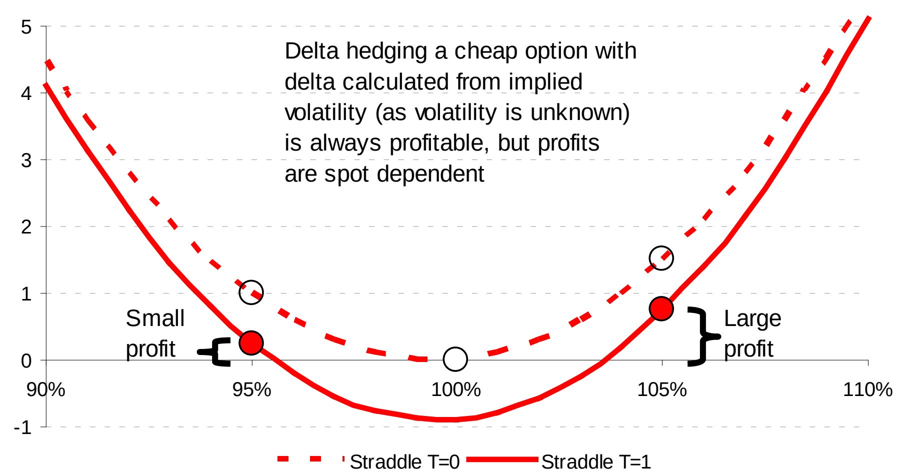

Source: Santander Investment Bolsa.

If implied volatility = realised volatility, profits are path independent

If the implied volatility is equal to the realised volatility, then the estimated delta calculated from the implied will be equal to the actual delta (calculated from the realised). In this case, profits from hedging will exactly match the theta cost for all paths, so it is path independent.

Figure 92. Profit (or Loss) from Continuously Delta Hedging Unknown Volatility

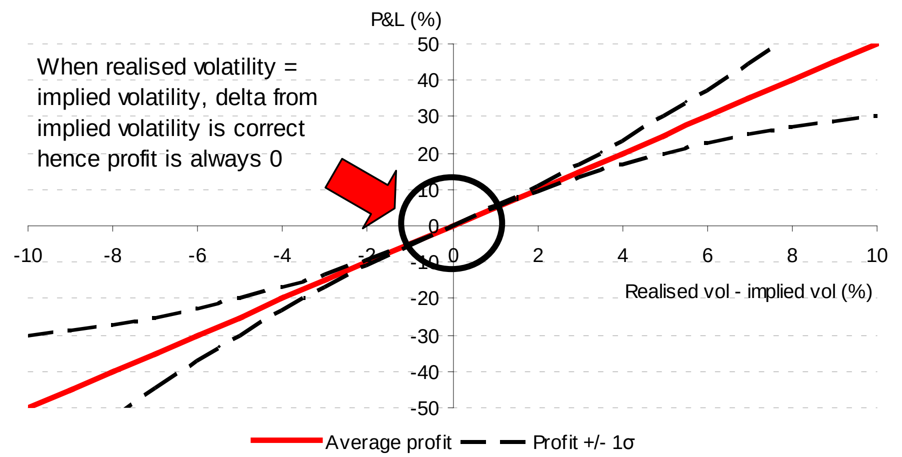

Source: Santander Investment Bolsa.

With continuous hedging, buying a cheap option is always profitable

If there is a difference between the actual delta and estimated delta, there is market risk but not enough to make a cheap option unprofitable (or an expensive option profitable). This is because in each infinitesimally small amount of time a cheap option will always reveal a profit from delta hedging (net of theta), although the magnitude of this profit is uncertain. The greater the difference between implied and realised, the greater the market risk and the larger the potential variation in profit.

DISCRETELY DELTA HEDGING WITH KNOWN VOLATILITY

While assuming continuous delta hedging is mathematically convenient, it is impossible in practice. Issues such as the cost of trading and minimum trading size (even if this is one share) make continuous trading impossible, as do fundamental reasons, such as trading hours (if you cannot trade 24 hours then it is impossible to trade overnight and prices can jump between the close of one day and start of another) and weekends.

Discrete hedging errors can be reduced by increasing the frequency of hedging

Discrete delta hedging adds noise to returns

The more frequent the discrete hedging, the less variation in the returns. If 24-hour trading were possible, then with an infinite frequency of hedging with known volatility the returns converge to the same case as continuous hedging with known volatility (ie, Black-Scholes). In order to show how the frequency of hedging can affect the payout of delta hedging, we shall examine hedging for every 5% and 10% move in spot.

Figure 93. Profit from Discrete Delta Hedging with Different Frequencies

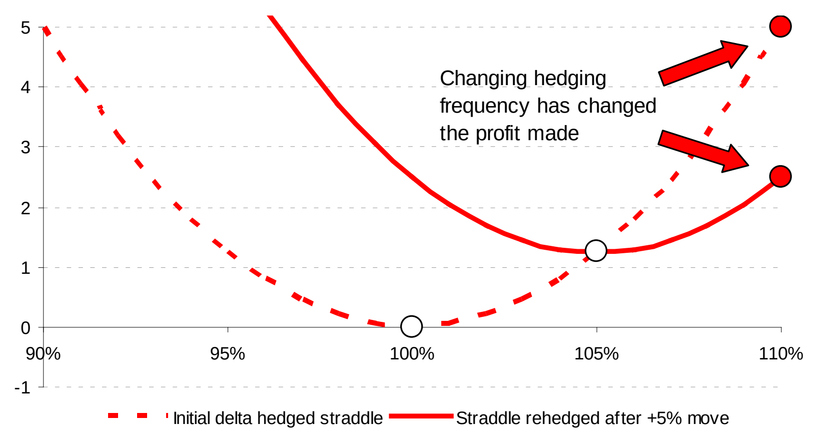

Source: Santander Investment Bolsa.

Hedging every 5% move in spot

If an investor delta hedges every 5% move in spot, then an identical profit is earned if the underlying rises 10% as if the underlying rises 5% and then returns to its starting point. This shows that the hedging frequency should ideally be frequent enough to capture the major turning points of an underlying.

Hedging every 10% move in spot

If the investor is hedged for every 10% move in the underlying, then no profit will be earned if the underlying rises 5% and then returns to its starting point. However, if the underlying rises 10%, a far larger profit will be earned than if the position was hedged every 5%. This shows that in trending markets it is more profitable to let positions run than to re-hedge them frequently.

Hedging error is independent of average profitability of trade

Noise from discrete delta hedging is independent of how cheap the option is

As the volatility of the underlying is known, there is no error due to the calculation of delta. As the only variation introduced is essentially 'noise', the size of this noise, or variation, is independent from the average profitability (or difference between realised vol and implied vol) of the trade.

Figure 94. Profit (or Loss) from Discrete Delta Hedging Known Volatility

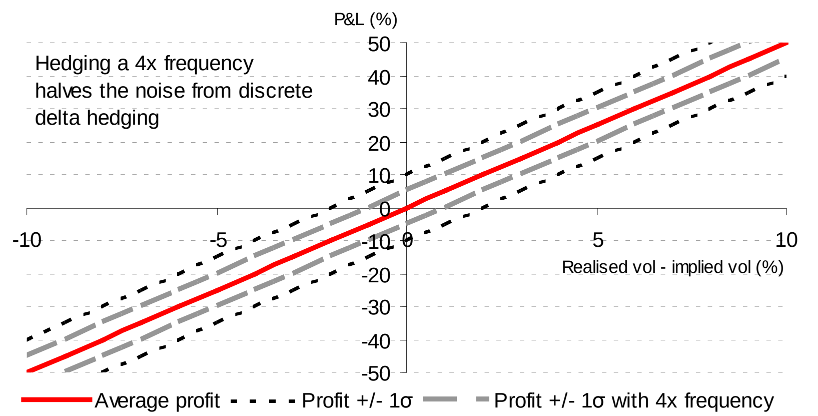

Source: Santander Investment Bolsa.

With discrete hedging, cheap options can lose money

With continuous delta hedging (with known or unknown volatility) it is impossible to lose money on a cheap option (an option whose implied volatility is less than the realised volatility over its life). However, as the error from discrete hedging is independent from the profitability of the trade, it is possible to lose money on a cheap option (and make money on an expensive option).

Hedging error is halved if frequency of hedging increased by factor of four

The size of the hedging error can be reduced by increasing the frequency of hedging. An approximation (shown below) is that if the frequency of hedging is increased by a factor of four, the hedging error term halves. This rule of thumb breaks down for very high-frequency hedging, as no frequency of hedging can eliminate the noise from non-24x7 trading (it will always have noise, due to the movement in share prices from one day's close to the next day's open).

$$\sigma_{P\&L} \approx \sigma \times vega \times \sqrt{\frac{\pi}{4N}}$$ where N is the number of times position is hedged in a year

DISCRETE DELTA HEDGING WITH UNKNOWN VOLATILITY

The most realistic assumption for profitability comes from the combination of discrete delta hedging and unknown volatility. Trading hours and trading costs are likely to limit the frequency at which a trader can delta hedge. Equally, the volatility of a stock is unknown, so implied volatility is likely to be used to calculate the delta. The variation in the profit (or loss) is caused by the variation due to discrete hedging and the inaccuracy of the delta (as volatility is unknown). Figure 95 below shows this combined variation in profit (or loss) and, as for discrete hedging with known volatility, it is possible for a delta hedged cheap option to reveal a loss.

Figure 95. Profit (or Loss) from Discrete Delta Hedging with Unknown Volatility

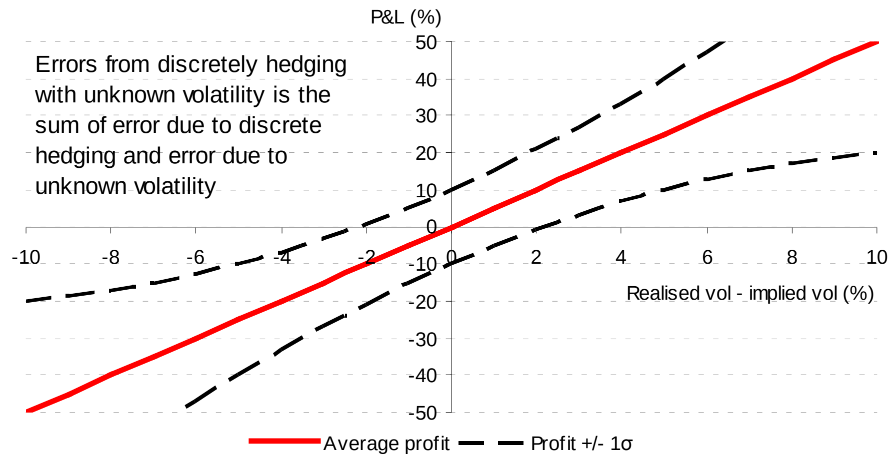

Source: Santander Investment Bolsa.

USE EXPECTED, NOT IMPLIED VOL FOR DELTA CALCULATIONS

As an example, let us assume a Dec08 SX5E ATM straddle was purchased in April 2008. In theory, it should be very profitable as the realised volatility of 39% was more than 50% above the 25% implied. However, most of the volatility came after the Lehman bankruptcy, which occurred towards the end of the option's life. If implied volatility was used to calculate the delta, then the time value would be assumed to be near zero. As equity markets had declined since April, the strike of the straddle would be above spot, hence we would have a delta ≈ -100% (the call would be OTM with a delta ≈ 0, while the put would be ITM with a delta ≈ 100%). To be delta-hedged, the investor would then buy 100% of the underlying per straddle. If the delta was calculated using the actual volatility (which was much higher), then the time value would be higher and the delta greater than -100% (eg, -85%). As the delta-hedged investor would have bought less than 100% of the underlying per straddle, this position outperformed hedging with implied volatility when the market fell after Lehman collapsed (as delta was lower, so less of the underlying was bought).

These results can be seen in Figure 96 below, which gives a clear example of why traders should hedge with the delta calculated from expected volatility rather than implied volatility. Because of the extreme volatility at the end of 2008, the two deltas differed at times by 24% (60% vs 84%).

Figure 96. Payout from Delta Hedging with Implied vs Realised Volatility

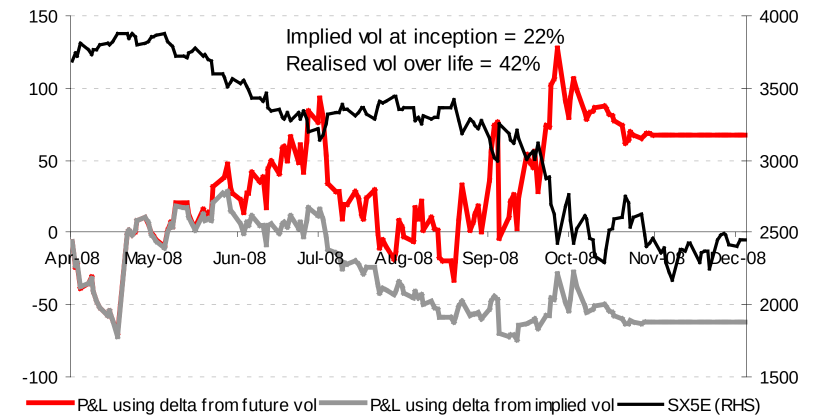

Source: Santander Investment Bolsa.

Hedging with delta using implied volatility is bad for long volatility strategies

Using expected, not implied, volatility to calculate delta is most important for long volatility strategies

Typically, when volatility rises there is often a decline in the markets. The strikes of the option are therefore likely to be above spot when actual volatility is above implied. This reduces the profits of the delta-hedged position as the position is actually long delta when it appears to be delta flat. Alternatively, the fact that the position hedged with the realised volatility over the life of the option is profitable can be thought of as due to the fact it is properly gamma hedged, as it has more time value than is being priced into the market. Hence, if a trader buys an option when the implied looks 5pts too cheap, then the hedge using delta should be calculated from a volatility 5pts above current implied volatility. Using the proper volatility means the profit is approximately the difference between the theoretical value of the option at inception (ie, using actual realised volatility in pricing) and the price of the option (ie, using implied volatility in pricing).

SKEW AND TERM STRUCTURE TRADING

SKEW AND TERM STRUCTURE ARE LINKED

When there is an equity market decline, there is normally a larger increase in ATM implied volatility at the near end of volatility surfaces than the far end. Assuming sticky strike, this causes near-dated skew to be larger than far-dated skew. The greater the term structure change for a given change in spot, the higher skew is. Skew is also positively correlated to term structure (this relationship can break down in panicked markets). For an index, skew (and potentially term structure) is also lifted by the implied correlation surface. Diverse indices tend to have higher skew for this reason, as the ATM correlation is lower (and low strike correlation tends to 100% for all indices).

SKEW AND TERM STRUCTURE CUT SURFACE IN DIFFERENT DIMENSIONS

A volatility surface has three dimensions (strike, expiry and implied volatility), which is difficult to show on a two dimensional page. For simplicity, a volatility surface is often plotted as two separate two dimensional graphs. The first plots implied volatility vs expiry (similar to the way in which a yield curve plots credit spread against expiry) in order to show term structure (the difference in implied volatility for options with different maturities and the same strike). The second plots implied volatility vs strike to show skew (the difference in implied volatility for options with different strikes and the same maturity). We examine a volatility surface in both these ways (ie, term structure and skew) and show how they are related.

TERM STRUCTURE IS NORMALLY UPWARD SLOPING

Near-dated implieds are more volatile than far-dated implieds

When there is a spike in realised volatility, near-dated implied volatility tends to spike in a similar way (unless the spike is due to a specific event such as earnings). This is because the high realised volatility is expected to continue in the short term. Realised volatility can be expected to mean revert over a c8-month period, on average. Hence far-dated implied volatilities tend to rise by a smaller amount than near-dated implied volatilities (as the increased volatility of the underlying will only last a fraction of the life of a far-dated option). Near-dated implieds are therefore more volatile than far-dated implieds. The theoretical term structure for different strikes is shown in Figure 97 below, which demonstrates that near-dated implieds are more volatile. We have shown ATM (100%) term structure as upward sloping as this is how it trades on average (for the same reasons credit spread term structure is normally upward sloping, ie, risk aversion and supply-demand imbalances for long maturities).

Figure 97. Term Structure for Options of Different Strikes

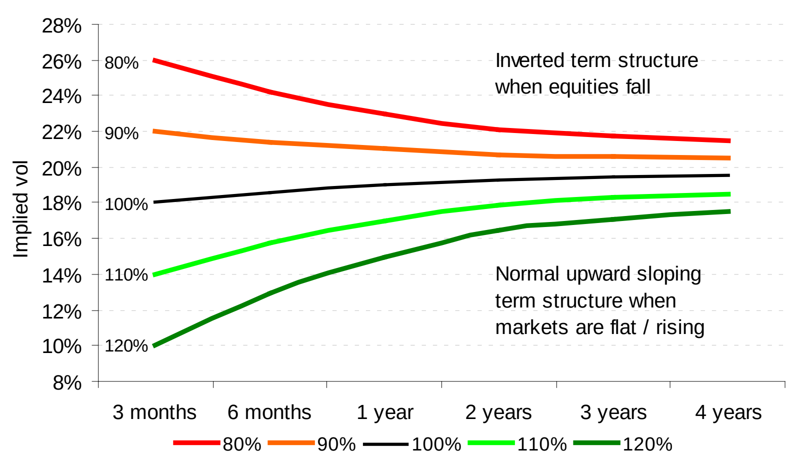

Source: Santander Investment Bolsa.

If equity markets decline, term structure becomes inverted

ATM term structure is typically positive in stable or rising markets, but negative in declining markets

Typically, an increase in volatility tends to be accompanied by a decline in equity markets, while a decline in volatility tends to occur in periods of calm or rising markets. If volatility surfaces are assumed not to move as spot moves (ie, sticky strike), then this explains why the term structure of low strike implied volatility is normally downward sloping (as the 80% strike term structure will be the ATM term structure when equities fall 20%). Similarly, this explains why the term structure of high strike implieds is normally upward sloping (as the 120% strike term structure will become the ATM term structure when equities rise 20%).

Slope of rising term structure is shallower than slope of inverted term structure

While Figure 97 above shows the term structure of a theoretical volatility surface, in practice the slope of rising term structure is shallower than the slope of inverted term structure. This can be seen by looking at a volatility cone (Figure 98). Despite the fact that the inverted term structure is steeper, the more frequent case of upward sloping term structure means the average term structure is slightly upward sloping[28].

Figure 98. Implied and Historic Volatility Cone (SX5E since 2006)

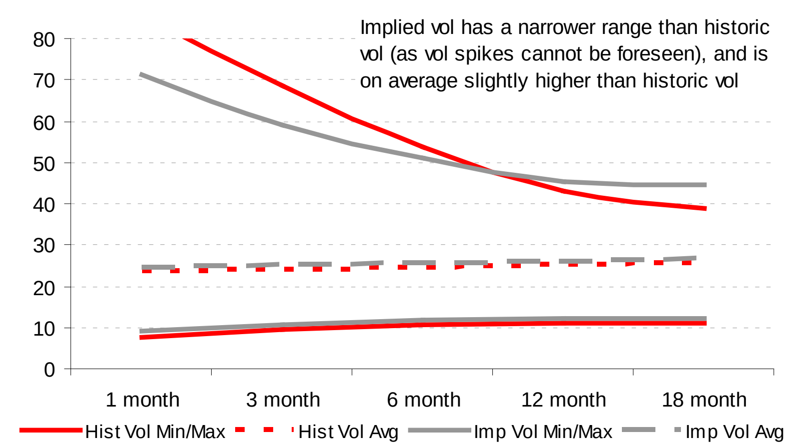

Source: Santander Investment Bolsa.

Implied volatility is usually greater than realised volatility and less volatile

While historic and realised volatility are linked, there are important differences which can be seen when looking at empirical volatility cones. Average implied volatility lies slightly above average realised volatility, as implieds are on average slightly expensive. Implied volatility is also less volatile (it has a smaller min-max range) than realised volatility for near-dated maturities. This is because implieds are forward looking (ie, similar to an average of possible outcomes) and there is never 100% probability of the maximum or minimum possible realised. This effect fades for longer maturities, potentially due to the additional volatility caused by supply-demand imbalances (eg, varying demand for put protection). This causes inverted implied volatility term structure to be less steep than realised volatility term structure.

[28] Positive implied correlation term structure will also lift index term structure relative to single stock.

SKEW IS INVERTED AND IS HIGHER FOR NEAR-DATED EXPIRIES

Near-term skew is steeper than far-dated skew as near-dated ATM is more volatile

Assuming volatility surfaces stay constant (ie, sticky strike), the effect of near-dated ATM implieds moving further than far-dated implieds for a given change in spot is priced into volatility surfaces by having a larger near-dated skew. The example data given in Figure 97 above is plotted in Figure 99 below with a change of axes to show skew for options of different maturity. This graph shows that near-dated implieds have higher skew than far-dated implieds. The more term structure changes for a given change in spot, the steeper skew is. As near-dated ATM volatility is more volatile than far-dated ATM volatility, near-dated implied volatility has higher skew.

Figure 99. Skew for Options of Different Maturity

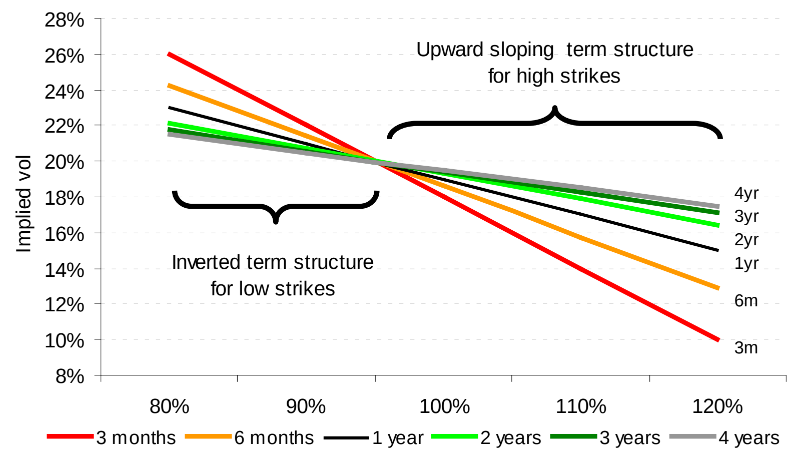

Source: Santander Investment Bolsa.

Skew for equities is normally inverted

Unless there is a high likelihood of a significant jump upwards (eg, if there were a potential takeover event), equities normally have negative skew (low strike implied greater than high strike implied). There are many possible explanations for this, some of which are listed below.

- **Big jumps in spot tend to be down, rather than up.** If there is a jump in the stock price, this is normally downwards as it is more common for an unexpectedly bad event to occur (bankruptcy, tsunami, terrorist attack, accident, loss or death of key personnel, etc) than an unexpectedly good event to occur (positive drivers are normally planned for).
- **Volatility is a measure of risk and leverage (hence risk) increases as equities decline.** If we assume no change in the number of shares in issue or amount of debt, then as a company's stock price declines its leverage (debt/equity) increases. Both leverage and volatility are a measure of risk and, hence, they are correlated, with volatility rising as equities fall.
- **Demand for protection and call overwriting.** Typically, investors are interested in buying puts for protection, rather than selling them. This lifts low strike implieds. Additionally, some investors like to call overwrite their positions, which weighs on higher strike implieds.

REASONS WHY SKEW AND TERM STRUCTURE ARE CORRELATED

The correlation between skew and term structure is shown below. The diagram only shows data for positive term structure, as the relationship tends to break down during a crisis.

Figure 100. SX5E Skew and Positive Term Structure (2007-10)[29]

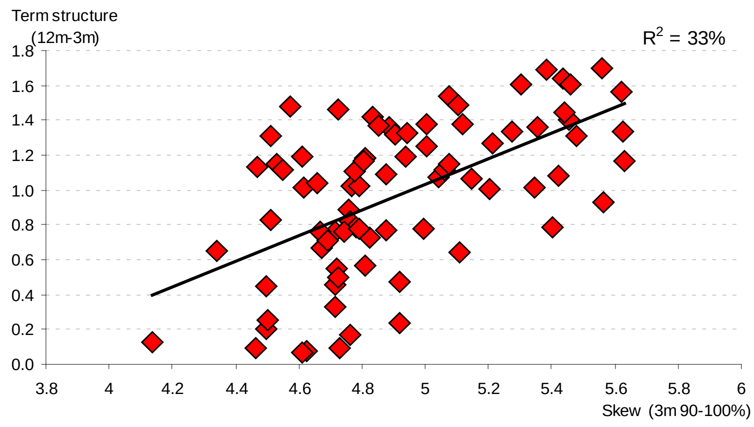

Source: Santander Investment Bolsa.

There are three reasons why skew and term structure are correlated

- Credit events, such as bankruptcy, lift both skew and term structure
- Implied volatility is 'sticky' for low strikes and long maturities
- Implied correlation is 'sticky' for low strikes and long maturities (only applies to index)

(1) BANKRUPTCY LIFTS BOTH SKEW AND TERM STRUCTURE

There are various models that show the effect of bankruptcy (or credit risk) lifting both skew and term structure. As implieds with lower strikes have a greater sensitivity to credit risk (as most of the value of low strike puts is due to the risk of bankruptcy), their implieds rise more, which causes higher skew. Similarly, options with longer maturity are more sensitive to credit risk (causing higher term structure, as far-dated implieds rise more). Longer-dated options have a higher sensitivity to credit risk as the probability of entering default increases with time (hence a greater proportion of an option's value will be associated with credit events as maturity increases). More detail on the link between volatility and credit can be seen in section *Capital Structure Arbitrage* in the *Appendix*.

[29] Excludes data from April 2010 onwards, as the change in US regulation regarding prop desks and moving equity derivatives onto exchanges (hence increased margin requirements) caused a spike in skew for major indices.

Figure 101. Term Structure Rising with Falling Volatility **Skew Rising with Falling Volatility**

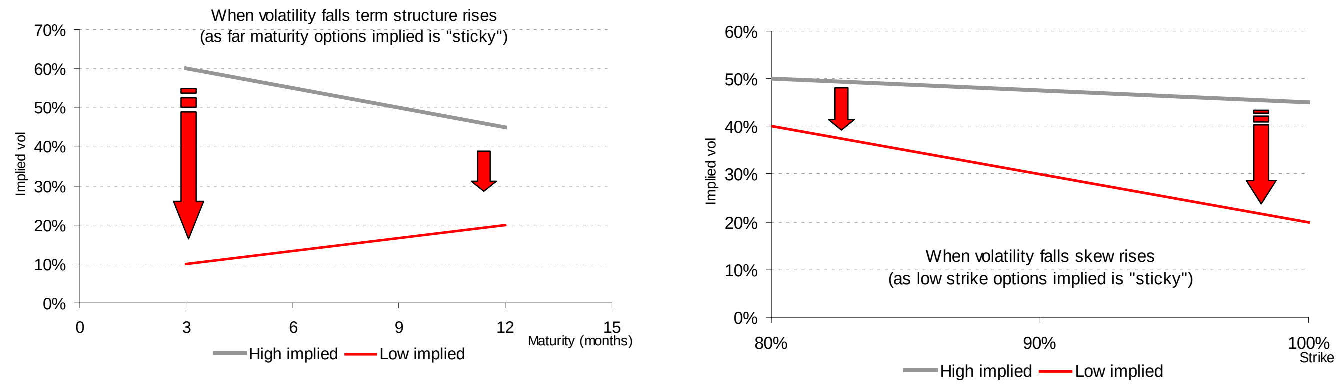

Source: Santander Investment Bolsa.

(2) IMPLIED VOL IS ‘STICKY’ FOR LOW STRIKES AND LONG MATURITIES

If there is a sudden decline in equity markets, it is reasonable to assume realised volatility will jump to a level in line with the peak of realised volatility. Therefore, low-strike, near-dated implieds should be relatively constant (as they should trade near the all-time highs of realised volatility). If a low-strike implied is constant, the difference between a low-strike implied and ATM implied increases as ATM implieds falls. This means near-dated skew should rise if near-dated ATM implieds decline (see Figure 101 above on the right). For this reason, we do not view skew as a reliable risk indicator, as it can be inversely correlated to ATM volatility[30]. The effect of falling implieds causing an increase in 90%-100% skew is shown with empirical data in Figure 102 below (we prefer to use 90%-100% skew rather than 90%-110%, as upside 100%-110% skew flattens as implieds reach a bottom).

Similarly, term structure should also rise if near-dated ATM implieds fall, as far-dated ATM implieds are relatively constant (as they tend to include complete economic cycles). This is shown in Figure 101 above on the left. Hence skew and term structure should be correlated as a fall in near-dated ATM implied lifts both of them.

Figure 102. SX5E 1 Year Max, Min and Average Implied Vol Since 2006

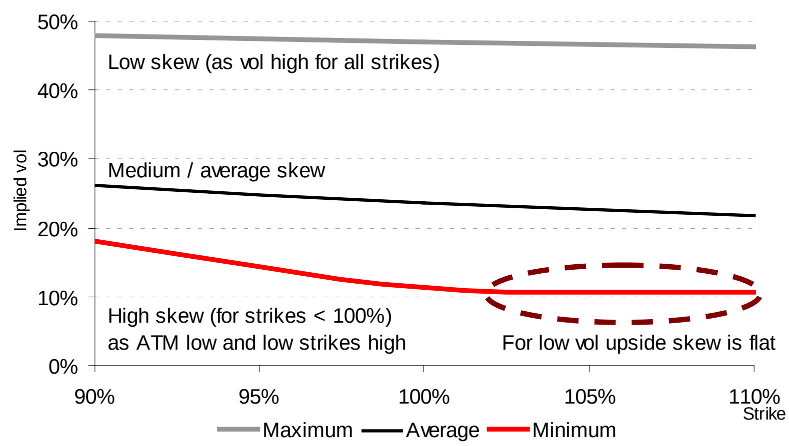

Source: Santander Investment Bolsa.

[30] ATM volatility is a risk measure; hence, a measure often inversely correlated to ATM volatility, such as skew, is not a reliable risk measure.

(3) CORRELATION SURFACE CAUSES INDEX SKEW AND TERM STRUCTURE TO BE CORRELATED

In the same way implied volatility is 'sticky' for low strikes and long maturities, so is implied correlation. This can be an additional reason why index skew and index term structure are correlated.

CORRELATION LIFTS INDEX SKEW ABOVE SINGLE-STOCK SKEW

Index skew is caused by both single-stock skew, and by the skew of the correlation surface

An approximation for implied correlation is the index volatility squared divided by the average single-stock volatility squared [$\rho = \sigma_{Index}^2 \div average(\sigma_{Single\ stock})^2$]. Implied correlation is assumed to tend towards 100% for low strikes, as all stocks can be expected to decline in a crisis. This causes index skew to be greater than single stock skew. Index skew can be thought of as being caused by both the skew of the single stock implied volatility surface, and the skew of the implied correlation surface.

Example of how index skew can be positive with flat single-stock skew

We shall assume all single stocks in an index have the same (flat) implied volatility and single-stock skew is flat. Low strike index volatility will be roughly equal to the constant single-stock volatility (as implied correlation is close to 100%), but ATM index volatility will be less than this figure due to diversity (as implied correlation ρ for ATM strikes is less than 100% and $\sigma_{Index}^2 = \rho \times average(\sigma_{Single\ stock})^2$. Despite single stocks having no skew, the index has a skew (as low strike index implieds > ATM index implieds) due to the change in correlation. For this reason, index skew is always greater than the average single-stock skew.

Implied correlation is likely to be sticky for low strikes and long maturities

A correlation surface can be constructed for options of all strikes and expiries, and this surface is likely to be close to 100% for very low strikes. The surface is likely to be relatively constant for far maturities; hence, implied correlation term structure and skew will be correlated (as both rise when near-dated ATM implied correlation falls, similar to volatility surfaces). This also causes skew and term structure to be correlated for indices.

DIVERSE INDICES HAVE HIGHER SKEW THAN LESS DIVERSE INDICES

As index skew is caused by both single-stock skew and implied correlation skew, a more diverse index should have a higher skew than a less diverse index (assuming there is no significant difference in the skew of the single-stock members). This is due to the fact that diverse indices have a lower ATM implied, but low strike implieds are in line with (higher) average single-stock implieds for both diverse and non-diverse indices.

SQUARE ROOT OF TIME RULE CAN COMPARE DIFFERENT TERM STRUCTURES AND SKEWS

When implied volatility changes, the change in ATM volatility multiplied by the square root of time is generally constant. This means that different (T_2-T_1) term structures can be compared when multiplied by $\sqrt{(T_2 - T_1)}/(\sqrt{T_2}-\sqrt{T_1})$, as this normalises against 1Y-3M term structure. Skew weighted by the square root of time should also be constant. Looking at the different term structures and skews, when normalised by the appropriate weighting, can allow us to identify calendar and skew trades in addition to highlighting which strike and expiry is the most attractive to buy (or sell).

REALISED VOLATILITY MEAN REVERTS AFTER EIGHT MONTHS

When there is a spike in realised volatility, it takes on average eight months for three-month realised volatility to settle back down to levels seen before the spike. The time taken for volatility to normalise is generally longer if the volatility is caused by a negative return, than if it is caused by a positive return (as a negative return is more likely to be associated with an event that increases uncertainty than a positive return). This mean reversion is often modelled via the square root of time rule.

VOLATILITY MOVE MULTIPLIED BY √TIME IS USUALLY CONSTANT

Square root of time rule is a quick way to sensibly adjust an entire volatility surface with just one parameter

The near-dated end of volatility surfaces is highly correlated to realised volatility, as hedge funds and prop desks typically initiate long/short gamma positions should there be a significant divergence. As volatility mean reverts, the far-dated end of volatility surfaces is more stable (as investors know that any spike in volatility will be short-lived and not last for the full length of a far-dated option). A common way to model the movement of volatility surfaces, is to define the movement of one-year implied and then adjust the rest of the curve by that move divided by time (in years) to the power of p. Only two parameters (the one-year move and p) are needed to adjust the whole surface. Fixing the power (or p) at 0.5 is the most common and is known as the square root of time rule (which only has one parameter, the one-year change).

$$\text{Implied volatility move for maturity T years} = \frac{\text{One year implied volatility move}}{T^{p}}$$

Figure 103. ATM Implied Volatility Moving in a Square Root of Time Manner

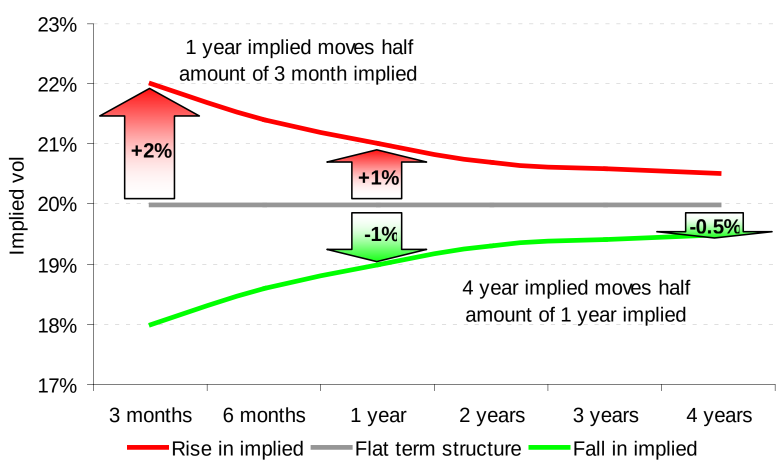

Source: Santander Investment Bolsa.

Square root of time rule has power 0.5, parallel moves are power 0

Historically, the power is 0.44, not 0.5 (but very close to square root of time)

While the above method is usually used with power 0.5 (square root of time rule), any power can be used. If there is a parallel movement in volatility surfaces (all maturities move the same amount), then a power of 0 should be used. In practice, implied volatility tends to move with power 0.44, suggesting that surfaces move primarily in a square root of time manner but at times also in parallel. If implieds rise (or decline) in a square root of time manner when equities decline (or rise), then this causes skew to decay by the square root of time as well (assuming sticky strike). This means that the skews of different maturities can be compared with each other by simply multiplying the skew by the square root of the maturity (see Figure 104 below).

Figure 104. Skew by Maturity (with same skew when multiplied by square root of time)

Maturity	3 Months	6 Months	1 Year	2 Years	3 Years	4 Years
Time (years)	0.25	0.5	1	2	3	4
Square root of time	**0.5**	**0.71**	**1**	**1.41**	**1.73**	**2**
90% implied	22.0%	21.4%	21.0%	20.7%	20.6%	20.5%
100% implied	18.0%	18.6%	19.0%	19.3%	19.4%	19.5%
Skew (per 10% move spot)	**4.0%**	**2.8%**	**2.0%**	**1.4%**	**1.2%**	**1.0%**
Skew × square root of time	**2.0%**	**2.0%**	**2.0%**	**2.0%**	**2.0%**	**2.0%**

Source: Santander Investment Bolsa.

POSITIVE PUT/CALL SPREADS IMPLY √TIME RULE FOR SKEW

Structures such as put spreads or call spreads, which can only have a positive payout, must have a cost associated with them, or investors would simply purchase an infinite amount of them for zero cost (or small profit) and enter a position which could never suffer a loss. This means that when the strike of a put is increased, its premium must rise too (intuitively correct, as the strike is the amount of money you receive when you 'put' the stock, so the higher the strike the better). Conversely, as the strike of a call increases, its premium must decrease. It can be shown that enforcing positive put/call spreads puts a cap/floor on skew, which decays by the square root of time. This provides mathematical support for the empirical evidence, suggesting implied volatility should normally move in a power weighted (by square root of time) manner[31]. For more details, see the section *Modelling Volatility Surfaces* in the *Appendix*.

Figure 105. SX5E Skew Multiplied by the Square Root of Time (R^2=83%)

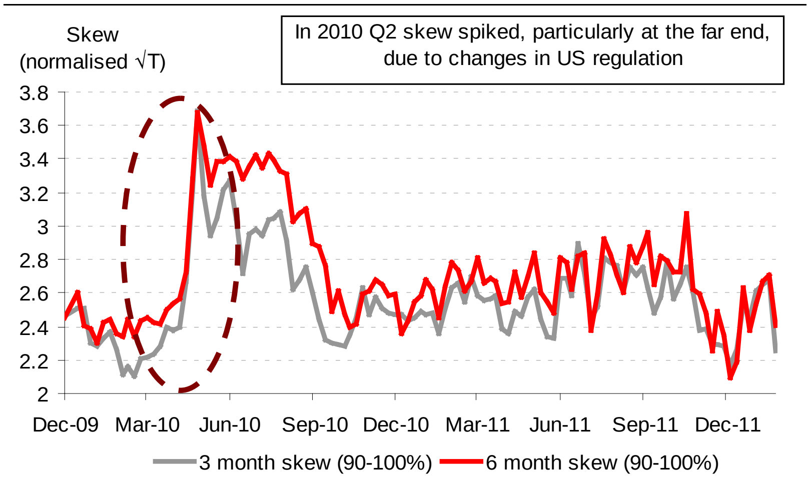

Source: Santander Investment Bolsa.

[31] Looking at ratio put spreads, it can be shown that for long maturities (five years) skew should decay by time, ie, 1/T or power=1 (rather than √T or power 0.5).

√TIME RULE CAN COMPARE DIFFERENT TERM STRUCTURES

ATM term structure can be modelled as flat volatility, with a square root of time adjustment on top. With this model, flat volatility is equal to the volatility for an option of infinite maturity. There are, therefore, two parameters to this model, the volatility at infinity (V_∞) and the scale of the square root of time adjustment ,which we define to be z (for one-year implied).

$$\text{Volatility} = V_\infty - z/\sqrt{T}$$

where

z = scale of the square root of time adjustment (which we define as normalised term structure)

We have a negative sign in front of z, so that a positive z implies an upward sloping term structure and a negative z is a downward sloping term structure.

Different term structures are normalised by multiplying by $\sqrt{(T_2T_1)}/(\sqrt{T_2}-\sqrt{T_1})$

Using the above definition, we can calculate the normalised term structure z from two volatility points V_1 and V_2 (whose maturity is T_1 and T_2).

$$V_1 = V_\infty - z/\sqrt{T_1}$$

$$V_2 = V_\infty - z/\sqrt{T_2}$$

⇨ $V_1 + z/\sqrt{T_1} = V_2 + z/\sqrt{T_2} = V_\infty$

⇨ $z\,(1/\sqrt{T_1} - 1/\sqrt{T_2}) = V_2 - V_1$

⇨ $z = (V_2 - V_1) \times \dfrac{\sqrt{T_2T_1}}{\sqrt{T_2} - \sqrt{T_1}}$

Normalising term structure by $\sqrt{(T_2T_1)}/(\sqrt{T_2}-\sqrt{T_1})$ puts it in the same 'units' as 1Y-3M term structure

$V_2 - V_1$ is the normal definition for term structure. Hence, term structure can be normalised by multiplying by $\sqrt{(T_2T_1)}/(\sqrt{T_2}-\sqrt{T_1})$. We note that the normalisation factor for 1Y-3M term structure is 1. Therefore, normalising allows all term structure to be compared to 1Y-3M term structure.

Figure 106. SX5E Normalised Term Structure (R2=80%)

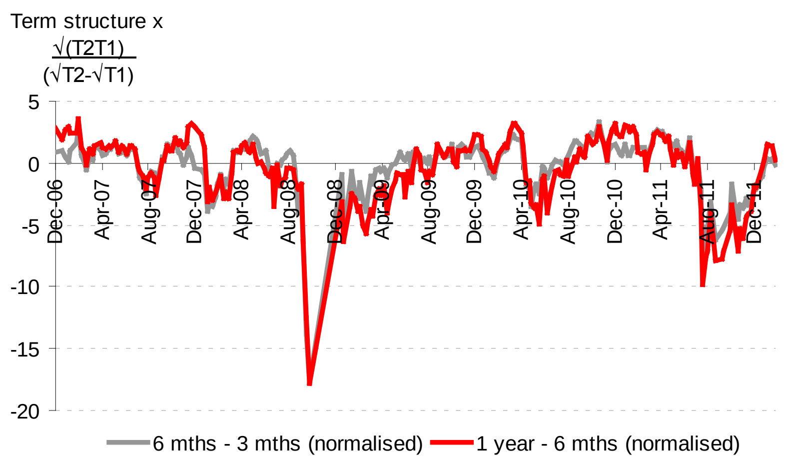

Source: Santander Investment Bolsa.

TERM STRUCTURE TRADES CAN PROFIT FROM IMBALANCES

The supply-demand imbalances of different products on implied volatility surfaces can create opportunity for other investors. The degree of the imbalance depends on the popularity of the product at the time. Investors who are willing to take the other side of the trade should be able to profit from the imbalance, and the risk taken can be hedged with other maturities or related securities.

Figure 107. Implied Volatility Imbalances by Maturity

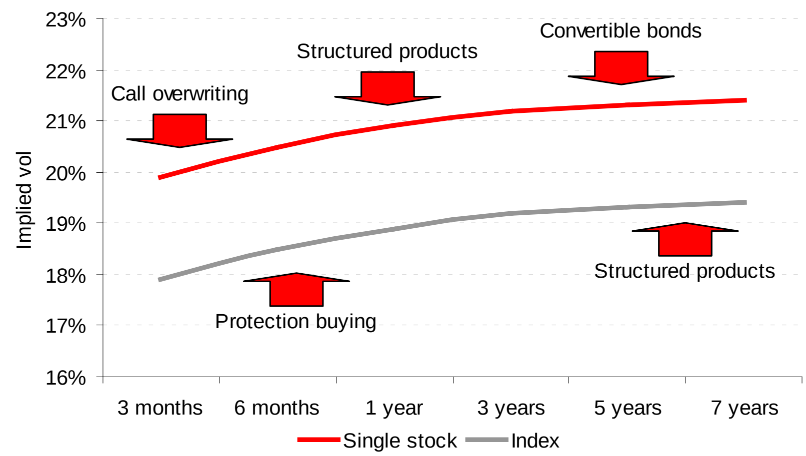

Source: Santander Investment Bolsa.

CALENDARS REMAIN CONSTANT IF SURFACES FOLLOW √TIME RULE

Calendars can be used to trade term structure imbalances

Given that the square root of time appears in the Black-Scholes formula for premium, the price of a 1x1 calendar (long one far-dated option, short one near-dated option) remains approximately constant if implied volatility surfaces move in a square root of time manner. Calendars can therefore be used to trade term structure imbalances as the trade is indifferent to the level of volatility as long as volatility moves in a power weighted manner.

IDENTIFYING WHEN TO GO LONG, OR SHORT, CALENDARS

When examining term structure trades, the power of the movement in volatility surfaces can be compared to the expected 0.5 power of the square root of time rule. If the movement has a power significantly different from 0.5, then a long (or short) calendar position could be initiated to profit from the anticipated correction. This method assumes calendars were previously fairly priced (otherwise the move could simply be a mean reversion to the norm).

If volatility rises with power less than 0.5, investors should short calendars

If surfaces rise with a power less than 0.5 (ie, a more parallel move) then near-dated implieds have not risen as much as expected and a short calendar (long near-dated, short far-dated) position should be initiated. This position will profit from the anticipated correction. Should surfaces fall with a power less than 0.5, a long calendar (short near-dated, long far-dated) would profit from the anticipated further decline of near-dated implieds.

If volatility rises with power more than 0.5, investors should go long calendars

Conversely, if surfaces rise with a power greater than 0.5, near-dated implieds have risen too far and a long calendar position should be initiated. On the other hand, if surfaces fall with a power greater than 0.5, a short calendar position should be initiated (as near-dated implieds have fallen too far).

POWER VEGA IS VEGA DIVIDED BY THE SQUARE ROOT OF TIME

As volatility surfaces tend to move in a square root of time manner, many systems report power vega (vega divided by square root of time). Power vega takes into account the fact that the implied volatility of near-dated options is more volatile than far-dated options.

VARIANCE TERM STRUCTURE CAN INDENTIFY TRADES

Variance term structure is similar to ATM term structure

To determine if a term structure trade is needed, we could look at variance term structure rather than implied volatility term structure. Using variance term structure eliminates the need to choose a strike (an ATM term structure will not be ATM as soon as the spot moves, so it is effectively strike dependent, but simply delegates the choice of strike to the equity market). Variance term structure is similar to ATM term structure, despite variance being long skew and skew being greater for near-dated implieds. This is because the time value of an OTM option increases with maturity. Hence, the increased weight associated with OTM options cancels the effect of smaller skew for longer maturities.

Forward starting variance swaps (or options) can be used to trade term structure

Trading term structure via a long and short variance swap is identical to a position in a forward starting variance swap (assuming the weights of the long and short variance swap are correct; if not, there will be a residual variance swap position left over). The correct weighting for long and short variance swaps to be identical to a forward starting variance swap is detailed in the section *Forward Starting Products*. If an investor wants to trade term structure, but does not want to have exposure to current volatility (ie, wants to have zero theta and gamma), then forward starting products (variance swaps or options) can be used. Note that while forward starting products have no exposure to current realised volatility, they do have exposure to future expectations of volatility (ie, implied volatility hence has positive vega).

HOW TO MEASURE SKEW AND SMILE

The implied volatilities for options of the same maturity, but of different strike, are different from each other for two reasons. Firstly, there is skew, which causes low strike implieds to be greater than high strike implieds due to the increased leverage and risk of bankruptcy. Secondly, there is smile (or convexity/kurtosis), when OTM options have a higher implied than ATM options. Together, skew and smile create the 'smirk' of volatility surfaces. We look at how skew and smile change by maturity in order to explain the shape of volatility surfaces both intuitively and mathematically. We also examine which measures of skew are best, and why.

MOMENTS DESCRIBE THE PROBABILITY DISTRIBUTION

Moments 1-4 describe forward, variance, skew and kurtosis

In order to explain skew and smile, we shall break down the probability distribution of log returns into moments. Moments can describe the probability distribution[32]. From the formula below we can see that the zero-th moment is 1 (as the sum of a probability distribution is 1, as the probability of all outcomes is 100%). The first moment is the expected value (ie, mean or forward) of the variable. The second, third and fourth moments are variance, skew and kurtosis, respectively (see table on the left below). For moments of two or greater it is usual to look at central moments, or moments about the mean (we cannot for the first moment as the first central moment would be 0). We shall normalise the central moment by dividing it by σ^n in order to get a dimensionless measure. The higher the moment, the greater the number of data points that are needed in order to get a reasonable estimate.

$$\text{Raw moment} = \mathrm{E}(X^k) = \int_{-\infty}^{\infty} x^k f(x)$$

$$\text{Normalised central moment} = \mathrm{E}((X-\mu)^k)/\sigma^k = \int_{-\infty}^{\infty} (x-\mu)^k f(x)/\sigma^k$$

where

$f(x)$ is the probability distribution function

Figure 108. Moments | **Related Option Position**

Moment	Name	Decay/Movement over Time	Maturity Where Dominates Surface	Related (Long) Position	Key Greek for Position	P&L Driver for Position
1†	Forward (expected price)	Random walk	NA	Stock/futures	Delta	Price
2	Variance (volatility²)	Mean reverts	Far-dated/all maturities	ATM options	Vega	Implied volatility
3	Skew	Decay square root of time	Medium-dated	Risk reversal (long low strikes & short high strikes)	Vanna	Skew (90-110%)
4	Kurtosis	Decay by time	Near-dated	Butterfly (long wings, short body)	Volga (gamma of volatility)	Vol of vol

(†) First raw moment (other moments are normalised central moments).
Source: Santander Investment Bolsa.

[32] The combination of all moments can perfectly explain any distribution as long as the distribution has a positive radius of convergence or is bounded (eg, a sine wave is not bounded; hence, it cannot be explained by moments alone).

VEGA MEASURES SIZE OF VOLATILITY POSITION

Vega measures the change in price of an option for a given (normally 1%) move in implied volatility. Implied volatility for far-dated options is relatively flat compared to near-dated, as both skew and kurtosis decay with maturity. Vega is highest for ATM options, as can be seen in the right hand chart in Figure 109 below.

Figure 109. Moment 2 = Variance | **Distribution for Constant Volatility** | **Vega is Size of Volatility Position**

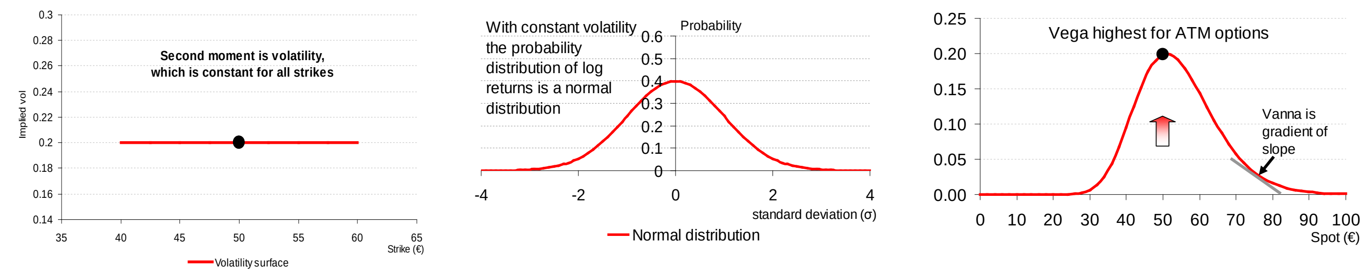

Source: Santander Investment Bolsa estimates.

VANNA MEASURES SIZE OF SKEW POSITION

Vanna (dVega/dSpot which is equal to dDelta/dVol) measures the size of a skew position[33], and is shown on the right side of Figure 110 below. Vanna is the slope of vega plotted against spot (see graph on right above).

Figure 110. Moment 3 = Skew[34] | **Distribution with Skew** | **Vanna is Size of Skew Position**

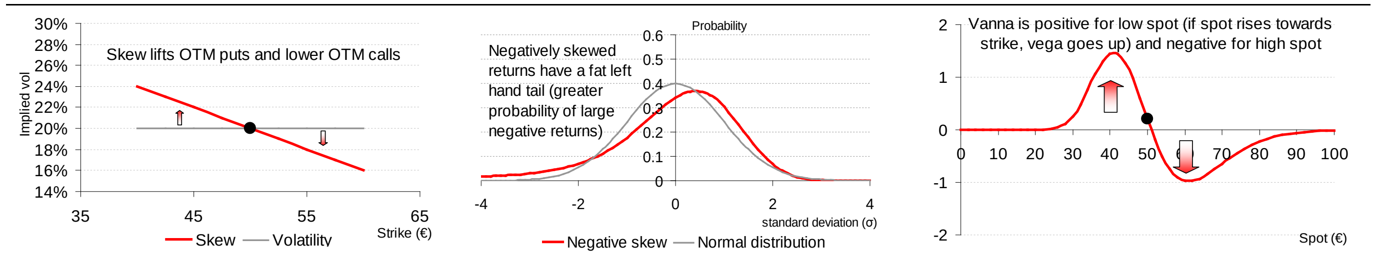

Source: Santander Investment Bolsa estimates.

VOLGA MEASURES SIZE OF GAMMA OF VOLATILITY

The gamma of volatility is measured by Volga (dVega/dVolatility), which is also known as volatility gamma or vega convexity. Volga is always positive (similar to option gamma always being positive) and peaks for c10-15 delta options (like Vanna).

Figure 111. Moment 4 = Kurtosis[34] | **Distribution with Kurtosis** | **Volga = Gamma of Vol**

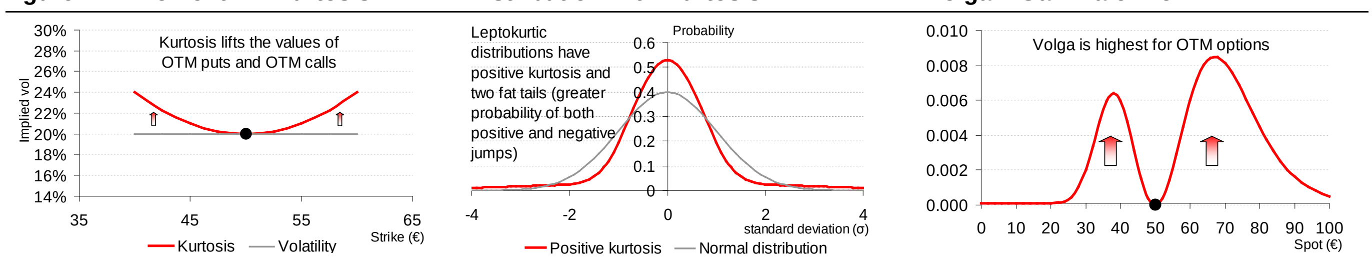

Source: Santander Investment Bolsa estimates.

[33] For details, see the next section *Skew Trading*.

[34] This is an approximation as the effect of moments on slope and convexity are intertwined.

Options with high volga benefit from volatility of volatility

Vol of vol gives volatility surfaces a 'smile' profile

Just as an option with high gamma benefits from high stock price volatility, an option with high volga benefits from volatility of volatility. The level of volatility of volatility can be calculated in a similar way to how volatility is calculated from stock prices (taking log returns is recommended for volatility as well). The more OTM an option is, the greater the volatility of volatility exposure. This is because the more implied volatility can change, the greater the chance of it rising and allowing an OTM option to become ITM. This gives the appearance of a 'smile', as the OTM option's implied volatility is lifted while the ATM implied volatility remains constant.

Stock returns have positive excess kurtosis and are leptokurtic

Kurtosis is always positive[35]. Hence, excess kurtosis (kurtosis -3) is usually used. The kurtosis (or normalised fourth moment) of the normal distribution is three; hence, normal distributions have zero excess kurtosis (and are known as mesokurtic). High kurtosis distributions (eg, stock price log returns) are known as leptokurtic, whereas low kurtosis distributions (pegged currencies that change infrequently by medium-sized adjustments) are known as platykurtic.

IMPLIED VOLATILITY SMIRK IS A COMBINATION OF SKEW AND SMILE

Kurtosis decays by 1/time, hence smile is most important for near-dated expiries

The final 'smirk' for options of the same maturity is the combination of skew (3^{rd} moment) and smile (4^{th} moment). The exact smirk depends on maturity. Kurtosis (or smile) can be assumed to decay with maturity by dividing by time[36] and, hence, is most important for near-dated expiries. For medium- (and long-) dated expiries, the skew effect will dominate kurtosis, as skew usually decays by the reciprocal of the square root of time (for more details, see the section *Modelling Volatility Surfaces* in the *Appendix*). Skew for equities is normally negative and therefore have mean < median < mode (max) and a greater probability of large negative returns (the reverse is true for positively skewed distributions). For far-dated maturities, the effect of both skew and kurtosis fades; hence, implieds converge to a flat line for all strikes. Skew can be thought of as the effect of changing volatility as spot moves, while smile can be thought of as the effect of jumps (up or down).

Figure 112. Near-Dated Implied Volatilities with Smirk (Skew and Smile)

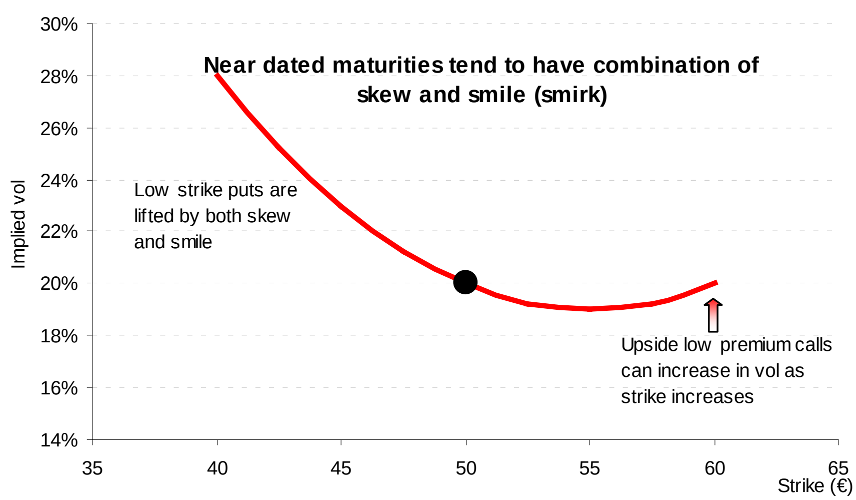

Source: Santander Investment Bolsa.

[35] Kurtosis is only zero for a point distribution.
[36] Assuming stock price is led by Lévy processes (eg, accumulation of independent identical shocks).

THERE ARE THREE WAYS TO MEASURE SKEW

There are three main ways skew can be measured. While the first is the most mathematical, in practice the other two are more popular with market participants.

- Third moment
- Strike skew (eg, 90%-110%)
- Delta skew (eg, [25 delta put – 25 delta call] / 50 delta)

(1) THIRD MOMENT IS DEFINITION OF CBOE SKEW INDEX

CBOE have created a skew index on the S&P500. This index is based on the normalised third central moment; hence, it is strike independent. The formula for the index is given below. For normal negative skew, if the size of skew increases, so does the index (as negative skew is multiplied by -10).

SKEW = 100 – 10 × 3rd moment

(2) STRIKE SKEW SHOULD NOT BE DIVIDED BY VOLATILITY

The most common method of measuring skew is to look at the difference in implied volatility between two strikes, for example 90%-110% skew or 90%-ATM skew. It is a common mistake to believe that strike skew should be divided by ATM volatility in order to take into account the fact that a 5pt difference is more significant for a stock with 20% volatility than 40% volatility. This ignores the fact that the strikes chosen (say 90%-110% for 20% volatility stocks) should also be wider for high volatility stocks (say 80%-120%, or two times wider, for 40% volatility stocks as the volatility is 2×20%). The difference in implied volatility should be taken between two strikes whose width between the strikes is proportional to the volatility (similar to taking the implied volatility of a fixed delta, eg, 25% delta). An approximation to this is to take the fixed strike skew, and multiply by volatility, as shown below. As the two effects cancel each other out, we can simply take a fixed strike skew without dividing by volatility.

Difference in vol between 2 strikes = 90-110%

⇨ Difference in vol between 2 strikes whose width increases with vol = 90-110% × ATM

$$\text{Skew} = \frac{\text{Difference in vol between 2 strikes whose width increases with vol}}{\text{ATM}}$$

$$\Rightarrow \text{Skew} = \frac{\text{90-110\%} \times \text{ATM}}{\text{ATM}}$$

⇨ Skew = 90-110%

Empirically, 90%-100% (or 90%-110%) skew is correct measure for fixed strike skew

The best measure of skew is one that is independent of the level of volatility. If this were not the case, then the measure would be partly based on volatility and partly on skew, which would make it more difficult to determine if skew was cheap or expensive. We have shown mathematically that an absolute difference (eg, 90%-110% or 90%-100%) is the correct measure of skew, but we can also show it empirically. The left-hand chart in Figure 113 below shows that there is no correlation between volatility and skew (90%-110%) for any European stocks that have liquid equity derivatives. If skew is divided by volatility, there is unsurprisingly a negative correlation between this measure and volatility (see right-hand chart below).

Figure 113. Strike Skew (90%-110%) Plotted vs Volatility **Strike Skew/Volatility Plotted vs Volatility**

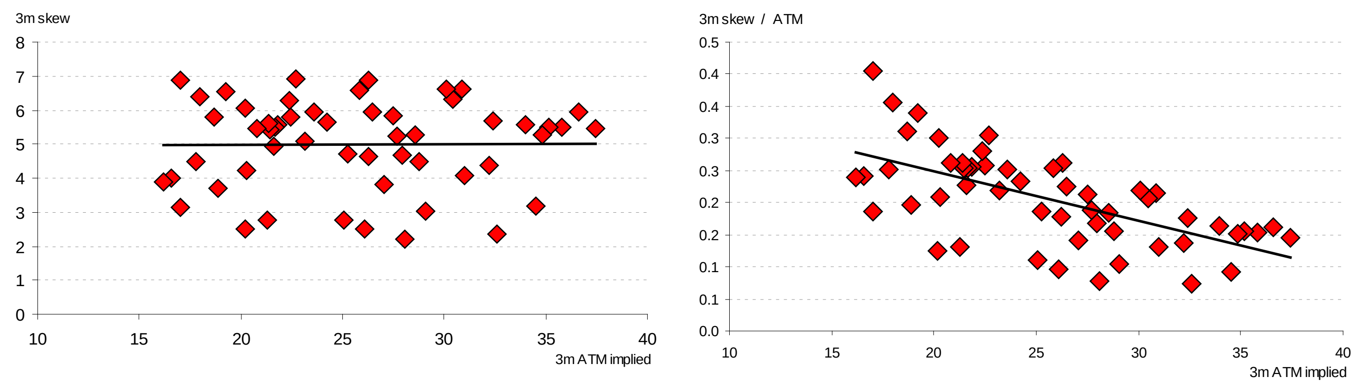

Source: Santander Investment Bolsa.

(3) DELTA SKEW IS VIRTUALLY IDENTICAL TO STRIKE SKEW

Arguably the best measure of skew is delta skew, where the difference between constant delta puts and calls is divided by 50 delta implied. An example of skew measured by delta is [25 delta put – 25 delta call] / 50 delta. As this measure widens the strikes examined as vol rises, in addition to normalising (ie, dividing) by the level of volatility, it is a 'pure' measure of skew (ie, not correlated to the level of volatility). While delta skew is theoretically the best measure, in practice it is virtually identical to strike skew. As there is a R^2 of 93% between delta skew and strike skew, we believe both are viable measures of skew (although strike skew is arguably more practical as it represents a more intuitive measure).

Figure 114. Strike Skew vs Delta Skew

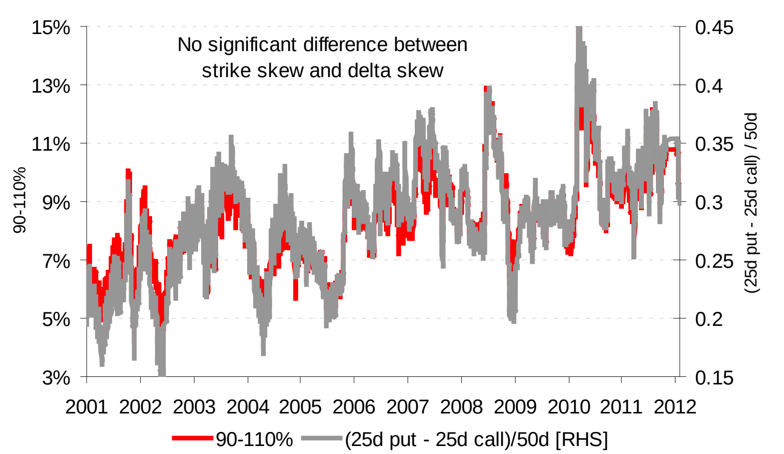

Source: Santander Investment Bolsa.

SKEW TRADING

The profitability of skew trades is determined by the dynamics of a volatility surface. We examine sticky delta (or 'moneyness'), sticky strike, sticky local volatility and jumpy volatility regimes. Long skew suffers a loss in both a sticky delta and sticky strike regimes due to the carry cost of skew. Long skew is only profitable with jumpy volatility. We also show how the best strikes for skew trading can be chosen.

FOUR IDEALISED REGIMES DESCRIBE MOVEMENT OF VOL SURFACE

There are four idealised regimes for a volatility surface. While sticky delta, sticky strike and (sticky) local volatility are well known and widely accepted names, we have added 'jumpy volatility' to define volatility with a high negative correlation with spot. These regimes are summarised below, and more details are given on pages 159-163 of this section.

(1) **Sticky delta (or sticky moneyness).** Sticky delta assumes a constant volatility for options of the same strike as a percentage of spot. For example, ATM or 100% strike volatility has constant volatility. As this model implies there is a positive correlation between volatility and spot, the opposite of what is usually seen in the market, it is not a particularly realistic model (except over a very long time horizon).

(2) **Sticky strike.** A sticky strike volatility surface has a constant volatility for options with the same fixed currency strike. Sticky strike is usually thought of as a stable (or unmoving) volatility surface as real-life options (with a fixed currency strike) do not change their implied volatility.

(3) **Sticky local volatility.** Local volatility is the instantaneous volatility of stock at a certain stock price. When local volatility is static, implied volatility rises when markets fall (ie, there is a negative correlation between stock prices and volatility). Of all the four volatility regimes, it is arguably the most realistic and fairly prices skew.

(4) **Jumpy volatility.** We define a jumpy volatility regime as one in which there is an excessive jump in implied volatility for a given movement in spot. There is a very high negative correlation between spot and volatility. This regime usually occurs over a very short time horizon in panicked markets (or a crash).

Figure 115. Volatility Surfaces When Equities Fall 10% **Volatility Surfaces When Equities Rise 10%**

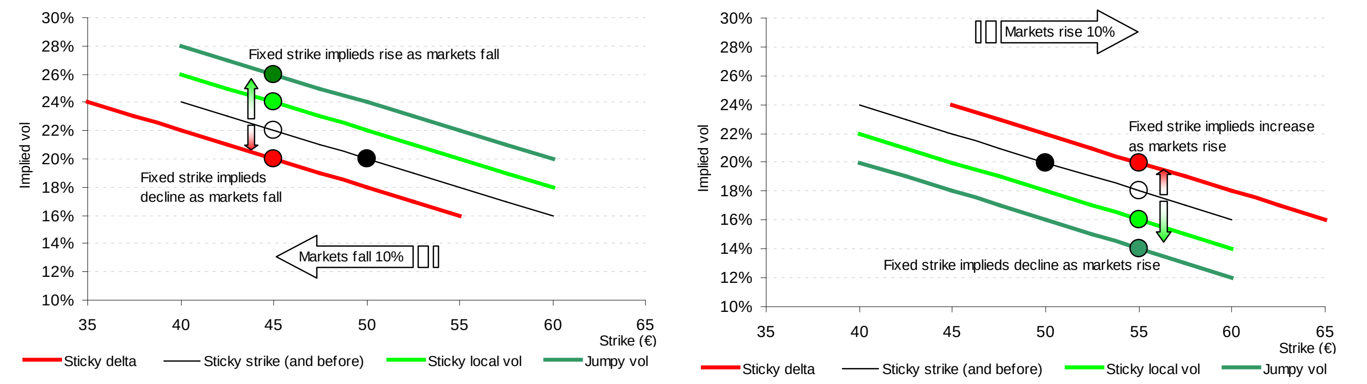

Source: Santander Investment Bolsa estimates.

PUTTING ON LONG SKEW TRADES HAS A COST (SKEW THETA)

Skew theta is the difference between the cost of gamma of an OTM option compared to an ATM option

If an investor initiates a long skew position by buying an OTM put and selling an OTM call, the implied volatility of the put purchased has a higher implied volatility than the implied volatility sold through the call. The long skew position therefore has a cost associated with it, which we shall define as 'skew theta'. Skew theta is the difference between the cost of gamma (theta per unit of dollar gamma) of an OTM option compared to an ATM option. If skew is flat, then all strikes have an identical cost of gamma, but as OTM puts have a higher implied volatility than ATM ones they pay more per unit of gamma. Skew theta is explained in greater depth at the end of this section. If the long skew position does not give the investor enough additional profit to compensate for the skew theta paid, then skew can be sold at a profit.

Skew trades profit from negative spot volatility correlation

Skew trades profit from negative spot volatility correlation

If there is a negative correlation between the movement of a volatility surface and spot (as is usually seen in practice), then this movement will give a long skew position a profit when the volatility surface is re-marked. For example, let us assume an investor is long skew via a risk reversal (long an OTM put and short an OTM call). If equity markets decline, the put becomes ATM and is the primary driver of value for the position (as the OTM call becomes further OTM it is far less significant). The rise in the volatility surface (due to negative correlation between spot and volatility) boosts the value of the (now ATM) put and, hence, the value of the risk reversal.

SKEW TRADES BREAK EVEN IF LOCAL VOL SURFACE IS CONSTANT

If the local volatility surface stays constant, the amount volatility surfaces move for a change in spot is equal to the skew (ie, ATM volatility moves by twice the skew, once for moving up the skew and another by the movement of the volatility surface itself). This movement is exactly the correct amount for the profit (or loss) on a volatility surface re-mark to compensate for the cost (or benefit) of skew theta[37]. The profit (or loss) caused by skew trades given the four volatility regimes are shown below.

Figure 116. Different Volatility Regimes and Breakdown of P&L for Skew Trades

	Fixed strike implied volatility change		P&L breakdown for long skew (e.g. long put, short call)				
Volatility regime	**Equity decline**	**Equity rise**	**Remark**		**Skew theta**		**Total**
Sticky delta	Falls	Rises		+		=	
Sticky strike	-	-		+		=	
Sticky local volatility	Rises	Falls		+		=	
Jumpy volatility	Rises significantly	Falls significantly		+		=	

Source: Santander Investment Bolsa.

[37] More details on local volatility can be found in the section *Local Volatility* in the *Appendix*.

SKEW IS USUALLY OVERPRICED DUE TO HEDGING

Demand for hedges causes skew to be overpriced

As volatility markets tend to trade between a static strike and static local volatility regime, long skew trades are usually unprofitable (usually there is negative spot volatility correlation, but not enough to compensate for the skew theta). As long skew trades break even during static local volatility regimes, they are only profitable in periods of jumpy volatility. This overpricing of skew can be considered to be a result of excessive demand for downside put options, potentially caused by hedging. Another reason for the overpricing of skew could be the popularity of short volatility long (downside) skew trades (traders often hedge a short volatility position with a long skew (OTM put) position, in order to protect themselves should markets suddenly decline). The profits from shorting expensive volatility are likely to more than compensate for paying an excessive amount for the long skew position (OTM put).

Figure 117. Market and Theoretical Skew

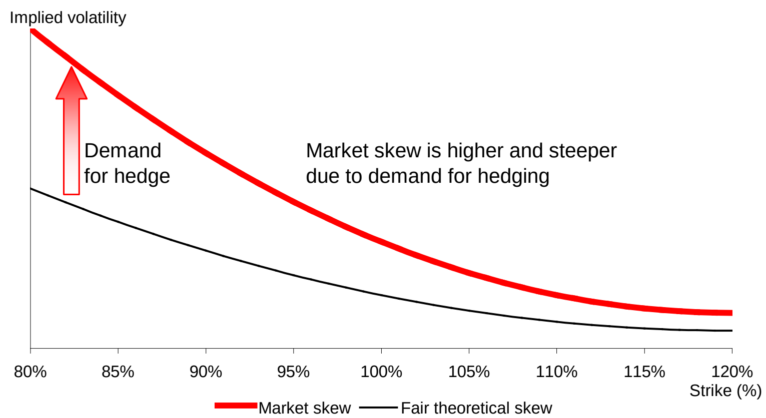

Source: Santander Investment Bolsa.

VOL REGIME DETERMINED BY TIME AND SENTIMENT

Implied volatility can be thought of as the equity derivative market's estimate of future volatility[38]. It is therefore investor sentiment that determines which implied volatility regime the market trades in, and this choice is largely determined by how much profit (or loss) a long skew position is expected to reveal over a certain time period (ie, investor sentiment). The choice of regime is also determined by the time horizon chosen.

Figure 118. Characteristics of Different Volatility Regimes

Characteristic	Sticky Delta		Sticky Strike		Sticky Local Vol		Jumpy Vol
Sentiment	Calm/trending		Normal				Panicked
Time horizon	Long term		Medium term				Short term
Spot vol correlation	Positive		Zero		Negative		Very negative
Call delta	δ_{call}	>	$\delta_{Black-Scholes}$	>	δ_{call}	>	δ_{call}
Put delta	δ_{put}	>	$\delta_{Black-Scholes}$	>	δ_{put}	>	δ_{put}
Abs(Put delta)	Abs(δ_{put})	<	Abs($\delta_{Black-Scholes}$)	<	Abs(δ_{put})	<	Abs(δ_{put})

Source: Santander Investment Bolsa estimates.

[38] In the absence of any supply-demand imbalance in the market.

Sticky delta regimes occur over long time horizon or trending markets

A sticky delta regime is typically one in which markets are trending in a stable manner (either up or down, with ATM volatility staying approximately constant) or over a very long time horizon of months or years (as over the long term the implied volatility mean reverts as it cannot go below zero or rise to infinity).

Jumpy volatility regimes occur over very short time horizons and panicked markets

It is rare to find a jumpy volatility regime that occurs over a long time horizon, as they tend to last for periods of only a few days or weeks. Markets tend to react in a jumpy volatility manner after a sudden and unexpected drop in equity markets (large increase in implied volatility given a decline in spot) or after a correction from such a decline (a bounce in the markets causing implied volatility to collapse).

Markets tend to trade between a sticky strike and sticky local volatility regime

On average, markets tend to trade between sticky strike and sticky local volatility

Sticky delta and jumpy volatility are the two extremes of volatility regimes. Sticky strike and sticky local volatility are far more common volatility regimes. Sticky strike is normally associated with calmer markets than sticky local volatility (as it is closer to a sticky delta model than jumpy volatility).

DELTA OF OPTION DEPENDS ON VOLATILITY REGIME

How a volatility surface reacts to a change in spot changes the value of the delta of the option. For sticky strike, as implied volatilities do not change, the delta is equal to the Black-Scholes delta.

Sticky delta regime implies a higher delta for options

However, if we assume a sticky delta volatility regime if an investor is long a call option, then the implied volatility of that option will decline if there is a fall in the market. The value of the call is therefore lower than expected for falls in the market. The reverse is also true as implied volatility increases if equities rise. As the value of the call is lower for declines and higher for rises (as volatility is positively correlated to spot), the delta is higher than that calculated by Black-Scholes (which is equal to the sticky strike delta).

A similar argument can be made for sticky local volatility (as volatility is negatively correlated to spot, the delta is less than that calculated by Black-Scholes). Figure 118 summarises the differences in delta for the different volatility regimes.

VOL CAN BE EXAMINED IN RELATIVE OR ABSOLUTE DIMENSIONS

To evaluate the profit – or loss – from a skew trade, assumptions have to be made regarding the movement of volatility surfaces over time (as we assume a skew trader always delta hedges, we are not concerned with the change in premium only the change in volatility). Typically, traders use two main ways to examine implied volatility surfaces. Absolute dimensions tend to be used when examining individual options, a snapshot of volatilities, or plotting implied volatilities over a relatively short period of time. Relative dimensions tend to be used when examining implied volatilities over relatively long periods of time[39].

- **Absolute dimensions.** In absolute dimensions, implied volatility surfaces are examined in terms of fixed maturity (eg, Dec14 expiry) and fixed strike (eg, €4,000). This surface is a useful way of examining how the implied volatility of actual traded options changes.
- **Relative dimensions.** An implied volatility surface is examined in terms of relative dimensions when it is given in terms of relative maturity (eg, three months or one year) and relative strike (eg, ATM, 90% or 110%). Volatility surfaces tend to move in relative dimensions over a very long period of time, whereas absolute dimensions are more suitable for shorter periods of time.

Care must be taken when examining implieds in relative dimensions

Relative dimensions should only be used for analysing long-term trends

As the options (and variance swaps) investors buy or sell are in fixed dimensions with fixed expiries and strikes, the change in implied volatility in absolute dimensions is the key driver of volatility profits (or losses). However, investors often use ATM volatility to determine when to enter (or exit) volatility positions, which can be misleading. For example, if there is a skew (downside implieds higher than ATM) and equity markets decline, ATM implieds will rise even though volatility surfaces remain stable. A plot of ATM implieds will imply buying volatility was profitable over the decline in equity markets; however, in practice this is not the case.

Absolute implied volatility is the key driver for equity derivative profits

As options that are traded have a fixed strike and expiry, it is absolute implied volatility that is the driver for equity derivative profits and skew trades. However, we accept that relative implied volatility is useful when looking at long-term trends. For the volatility regimes (1) sticky delta and (2) sticky strike, we shall plot implieds using both absolute and relative dimensions in order to explain the difference. For the remaining two volatility regimes (sticky local volatility and jumpy volatility), we shall only plot implied volatility using absolute dimensions (as that is the driver of profits for traded options and variance swaps).

[39] This is usually for liquidity reasons, as options tend to be less liquid for maturities greater than two years (making implied volatility plots of more than two years problematic in absolute dimensions).

(1) STICKY DELTA ASSUMES ATM VOL NEVER CHANGES

A sticky delta model assumes a constant implied volatility for strikes as a percentage of spot (eg, ATM stays constant). How a volatility surface moves with a change in spot is shown below for both absolute/fixed strike (Figure 119 on the left) and relative strike (Figure 119 on the right).

Figure 119. Sticky Delta Absolute/Fixed Strike **Sticky Delta Relative Strike (as Percentage of Spot)**

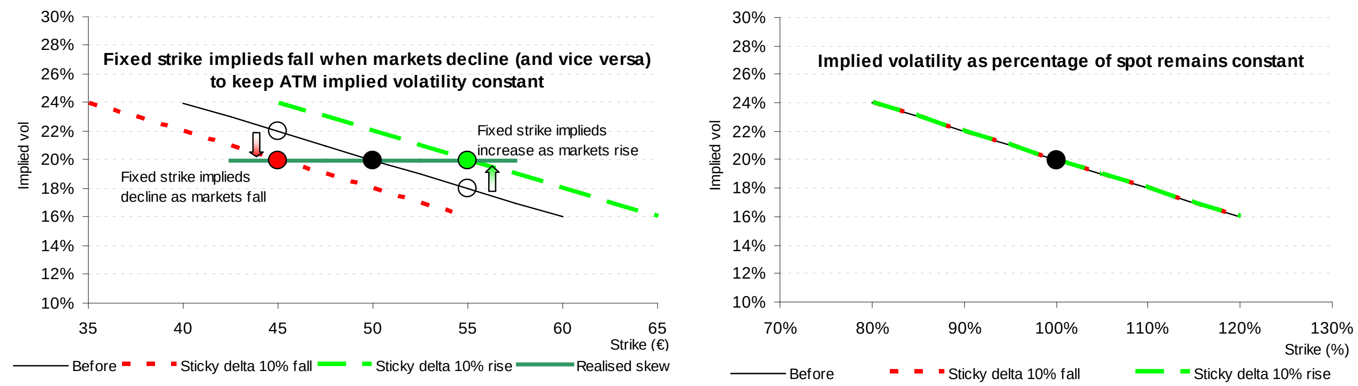

Source: Santander Investment Bolsa estimates.

RANGE-BOUND VOLATILITIES SUPPORT A STICKY DELTA MODEL

As implied volatility cannot be negative, it is therefore usually floored close to the lowest levels of realised volatility. Although an infinite volatility is theoretically possible, in practice implied volatility is typically capped close to the all-time highs of realised volatility. Over a long period of time, ATM implied volatility can be thought of as being range bound and likely to trend towards an average value (although this average value will change over time as the macro environment varies). As the trend towards this average value is independent of spot, the implied volatility surface in absolute dimensions (fixed currency strike) has to move to keep implied volatility surface in relative dimensions (strike as percentage of spot) constant. Thinking of implied volatility in this way is a sticky delta (or sticky moneyness) implied volatility surface model.

Sticky delta most appropriate over long term (many months or years)

Sticky delta most appropriate for time horizon of one year or more

While over the long term implied volatility tends to return to an average value, in the short term volatility can trade away from this value for a significant period of time. Typically, when there is a spike in volatility it takes a few months for volatility to revert back to more normal levels. This suggests a sticky delta model is most appropriate for examining implied volatilities for periods of time of a year or more. As a sticky delta model implies a positive correlation between (fixed strike) implied volatility and spot, the opposite of what is normally seen, it is not usually a realistic model for short periods of time. Trending markets (calmly rising or declining) are usually the only situation when a sticky delta model is appropriate for short periods of time. In this case, the volatility surface tends to reset to keep ATM volatility constant, as this implied volatility level is in line with the realised volatility of the market.

LONG SKEW IS UNPROFITABLE IN STICKY DELTA VOLATILITY REGIME

In a sticky delta volatility regime the fixed strike implied volatility (and, therefore, the implied volatility of traded options) has to be re-marked when spot moves. The direction of this re-mark for long skew positions causes a loss, as skew should be flat if ATM volatility is going to remain unchanged as markets move (we assume the investor has bought skew at a worse level than flat). Additionally, the long skew position carries the additional cost of skew theta, the combination of which causes long skew positions to be very unprofitable.

(2) STICKY STRIKE HAS ZERO SPOT VOL CORRELATION

A sticky strike model assumes that options of a fixed currency strike are fixed (absolute dimensions). The diagrams below show how a volatility surface moves in both absolute/fixed strike and relative strike due to a change in spot.

Figure 120. Sticky Strike Absolute/Fixed Strike **Sticky Strike Relative Strike (as Percentage of Spot)**

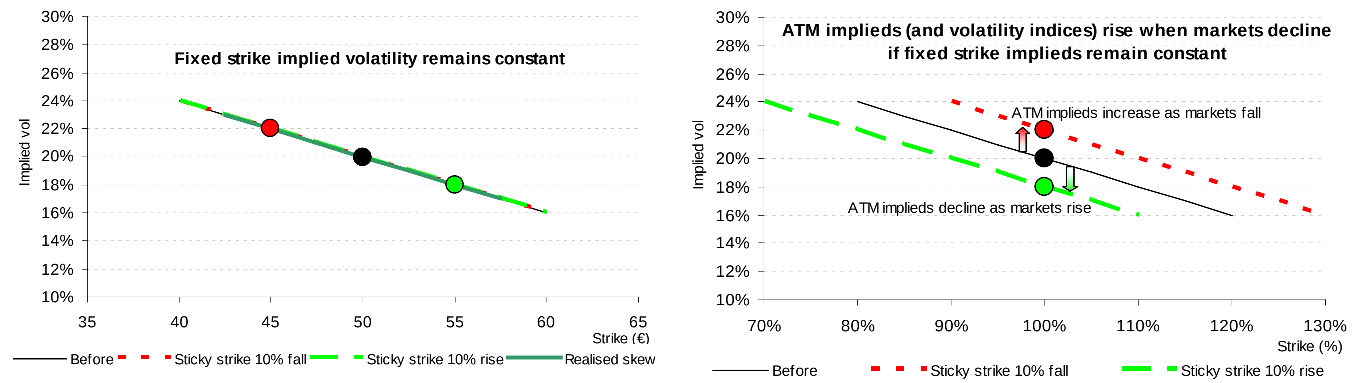

Source: Santander Investment Bolsa estimates.

TRADER'S SYSTEMS CAN GIVE ILLUSION OF STICKY STRIKE

While Figure 120 above describes which volatility regime normally applies in any given environment, there are many exceptions. A particular exception is that for very small time horizons volatility surfaces can seem to trade in a sticky strike regime. We believe this is due to many trading systems assuming a static strike volatility surface, which then has to be re-marked by traders (especially for less liquid instruments, as risk managers are likely to insist on volatilities being marked to their last known traded implied volatility)[40]. As the effect of these trading systems on pricing is either an illusion (as traders will re-mark their surface when asked to provide a firm quote) or well within the bid-offer arbitrage channel, we believe this effect should be ignored.

LONG SKEW UNPROFITABLE WITH STICKY STRIKE

Sticky strike can be thought of as a Black-Scholes model

While there is no profit or loss from re-marking a surface in a sticky strike model, a long skew position still has to pay skew theta. Overall, a long skew position is still unprofitable in sticky strike regimes, but it is less unprofitable than for a sticky delta regime.

[40] Anchor delta measures the effect of re-marking a volatility surface and is described in the section *Advanced (Practical or Shadow) Greeks* in the *Appendix*.

(3) STICKY LOCAL VOLATILITY PRICES SKEW FAIRLY

As a sticky local volatility causes a negative correlation between spot and Black-Scholes volatility (shown below), this re-mark is profitable for long skew positions. As the value of this re-mark is exactly equal to the cost of skew theta, skew trades break even in a sticky local volatility regime. If volatility surfaces move as predicted by sticky local volatility, then skew is priced fairly (as skew trades do not make a loss or profit).

BLACK-SCHOLES VOL IS AVERAGE OF INSTANTANEOUS LOCAL VOL

Local volatility is the name given for the instantaneous volatility of an underlying (ie, the exact volatility it has at a certain point). The Black-Scholes volatility of an option with strike K is equal to the average local (or instantaneous) volatility of all possible paths of the underlying from spot to strike K. This can be approximated by the average of the local volatility at spot and the local volatility at strike K. This approximation gives two results[41]:

- The ATM Black-Scholes volatility is equal to the ATM local volatility.
- Black-Scholes skew is half the local volatility skew (due to averaging).

Example of local volatility skew = 2x Black-Scholes skew

The second point can be seen if we assume the local volatility for the 90% strike is 22% and the ATM local volatility is 20%. The 90%-100% local volatility skew is therefore 2%. As the Black-Scholes 90% strike option will have an implied volatility of 21% (the average of 22% and 20%), it has a 90%-100% skew of 1% (as the ATM Black-Scholes volatility is equal to the 20% ATM local volatility).

Figure 121. Sticky Local Volatility Absolute/Fixed Strike

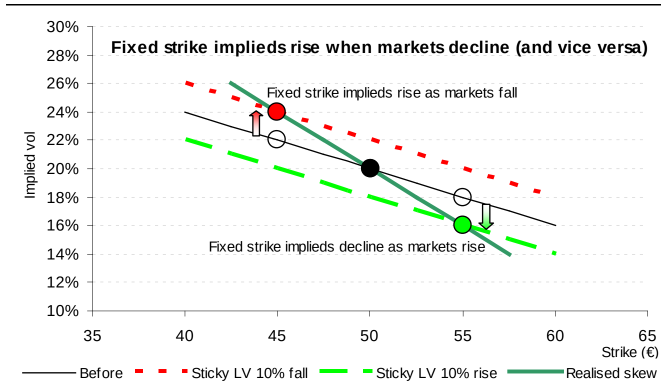

Black-Scholes and Local Volatility Skew

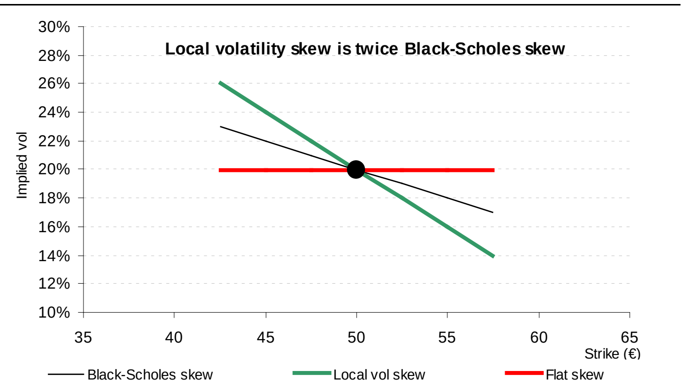

Source: Santander Investment Bolsa estimates.

STICKY LOCAL VOL IMPLIES NEGATIVE SPOT VOL CORRELATION

As local volatility skew is twice the Black-Scholes skew, and ATM volatilities are the same, a sticky local volatility surface implies a negative correlation between spot and implied volatility. This can be seen by the ATM Black-Scholes volatility resetting higher if spot declines and is shown in the diagrams above.

[41] As this is an approximation, there is a slight difference which we shall ignore.

Example of negative correlation between spot and Black-Scholes volatility

We shall use the values from the previous example, with the local volatility for the 90% strike = 22%, Black-Scholes of the 90% strike = 21% and the ATM volatility for both (local and Black-Scholes) = 20%. If markets decline 10%, then the 90% strike option Black-Scholes volatility will rise 1% from 21% to 22% (as ATM for both local and Black-Scholes volatility must be equal). This 1% move will occur in parallel over the entire surface (as the Black-Scholes skew has not changed). Similarly, should markets rise 10%, the Black-Scholes volatility surface will fall 1% (assuming constant skew).

LONG SKEW PROFITS FROM VOLATILITY SURFACE RE-MARK

In order to demonstrate how the negative correlation between spot and (Black-Scholes) implied volatility causes long skew positions to profit from volatility surfaces re-mark, we shall assume an investor is long a risk reversal (long OTM put, short OTM call). This position is shown in Figure 122 below. When markets fall, the primary driver of the risk reversal's value is the put (which is now more ATM than the call), and the put value will increase due to the rise in implied volatility (due to negative correlation with spot). Similarly, the theoretical value of the risk reversal will rise (as the call is now more ATM – and therefore the primary driver of value – and, as implied volatilities decline as markets rise, the value of the short call will rise as well). The long skew position therefore profits from both a movement up or down in equity markets, as can be seen in the diagram below as both the long call and short put position increase in value.

Figure 122. Premium of Long Put, Short Call (long skew) Risk Reversal

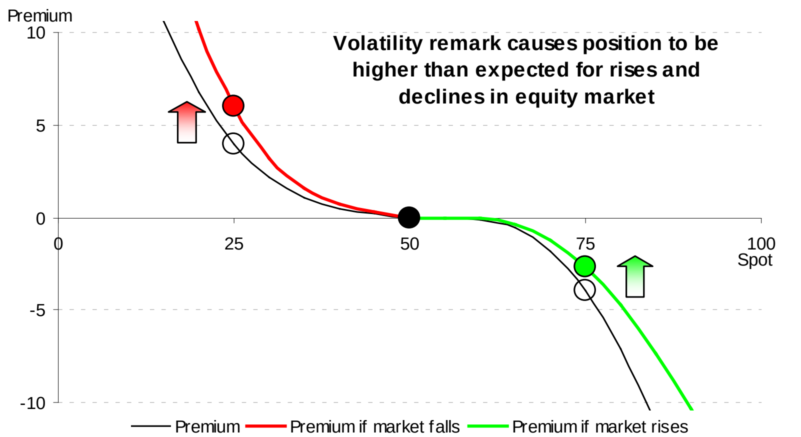

Source: Santander Investment Bolsa.

VOLATILITY RE-MARK WITH STICKY LOCAL VOL = SKEW THETA

Skew trades break even with a static local volatility model

While a sticky local volatility regime causes long skew positions to profit from (Black-Scholes) implied volatility changes, the position still suffers from skew theta. The combination of these two cancel exactly, causing a long (or short) skew trade to break even. As skew trades break even under a static local volatility model, and as there is a negative spot vol correlation, it is arguably the most realistic volatility model.

(4) JUMPY VOLATILITY IS ONLY REGIME WHERE LONG SKEW IS PROFITABLE

During very panicked markets, or immediately after a crash, there is typically a very high correlation between spot and volatility. During this volatility regime (which we define as jumpy volatility) volatility surfaces move in excess of that implied by sticky local volatility. As the implied volatility surface re-mark for a long skew position is in excess of skew theta, long skew positions are profitable. A jumpy volatility regime tends to last for a relatively short period of time.

Figure 123. Jumpy Volatility Absolute/Fixed Strike

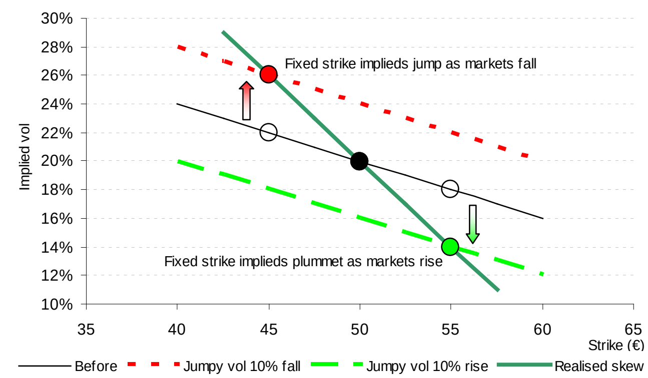

Realised Skew of Sticky Delta/Strike/Local/Jumpy Vol

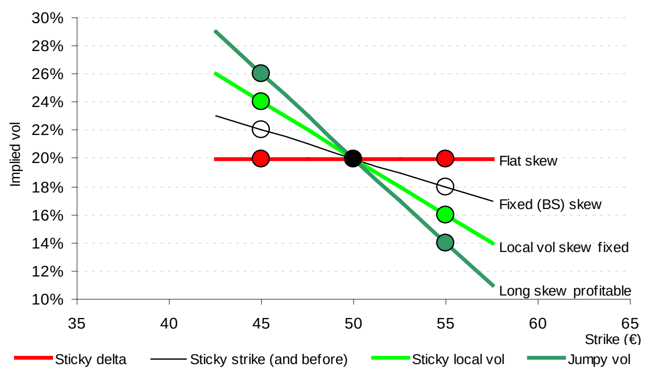

Source: Santander Investment Bolsa estimates.

Example of volatility regimes and skew trading

If one-year 90%-100% skew is 25bp per 1% (ie, 2.5% for 90-100%) and markets fall 1%, volatility surfaces have to rise by 25bp for the profit from realised skew to compensate for the cost of skew theta. If surfaces move by more than 25bp, surfaces are moving in a jumpy volatility way and skew trades are profitable. If surfaces move by less than 25bp then skew trades suffer a loss.

Figure 124. Breakdown of P&L for Skew Trades

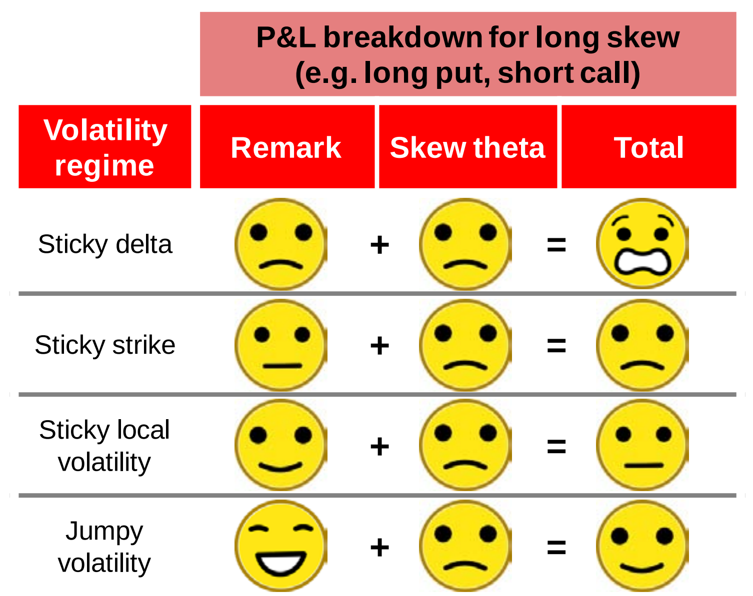

Source: Santander Investment Bolsa.

SKEW ONLY FAIRLY PRICED IF ATM MOVES BY TWICE THE SKEW

We define realised skew to be the profit due to re-marking the surface

For a given movement in spot from S_0 to S_1, we shall define the movement of the (Black-Scholes) implied volatility surface divided by the skew (implied volatility of strike S_1 – implied volatility of strike S_0) to be the realised skew. The realised skew can be thought of as the profit due to re-marking the volatility surface. Defining realised skew to be the movement in the volatility surface is similar to the definition of realised volatility, which is the movement in spot.

realised skew = movement of surface/skew

where:

movement of surface = movement of surface when spot moves from S_0 to S_1

skew = difference in implied volatility between S_1 and S_0

The ATM volatility can then be determined by the below equation:

$$ATM_{time\ 1} = ATM_{time\ 0} + skew + movement\ of\ surface$$

$$\Rightarrow ATM_{time\ 1} = ATM_{time\ 0} + skew + (skew \times realised\ skew)$$

$$\Rightarrow ATM_{time\ 1} = ATM_{time\ 0} + skew \times (1 + realised\ skew)$$

The realised skew for sticky delta is therefore -1 in order to keep ATM constant (and hence skew flat) for all movements in spot. A sticky strike regime has a realised skew of 0, as there is no movement of the volatility surface and skew is fixed. A local volatility model has a realised skew of 1, which causes ATM to move by twice the value implied by a fixed skew. As local volatility prices skew fairly, skew is only fairly priced if ATM moves by twice the skew. We shall assume the volatility surface for jumpy volatility moves more than it does for sticky local volatility, hence has a realised skew of more than 1.

Skew profit is proportional to realised skew – 1 (due to skew theta)

In order to calculate the relative profit (or loss) of trading skew, the value of skew theta needs to be taken away, and this value can be thought of as -1. Skew profit is then given by the formula below:

Skew profit α realised skew - 1

SKEW TRADING IS EQUIVALENT TO TRADING 2ND ORDER GAMMA

Comparing vanna to skew theta can identify trading opportunities

Determining the current volatility regime helps a trader decide if skew trades are likely to be profitable. In order to determine the strikes used to initiate long or short skew positions, a trader needs to evaluate the richness or cheapness of skew across different strikes. It is possible to show intuitively, and mathematically, that skew trading is very similar to delta hedging gamma. Given this relationship, comparing vanna (dVega/dSpot), weighted by the square root of time, to skew theta can be a useful rule of thumb to identify potential trading opportunities.

Figure 125. Call Option with 50 Strike

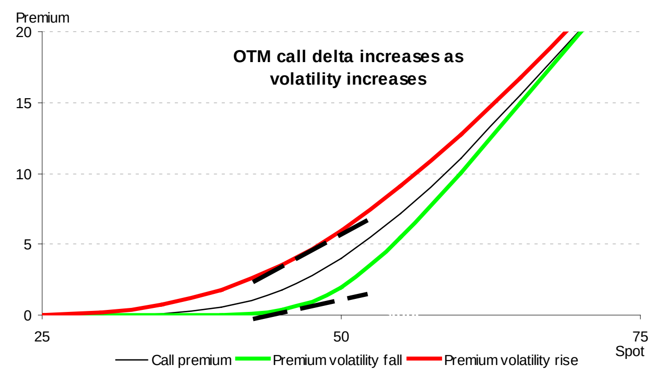

Delta of Call Option with Rise and Fall in Implied Volatility

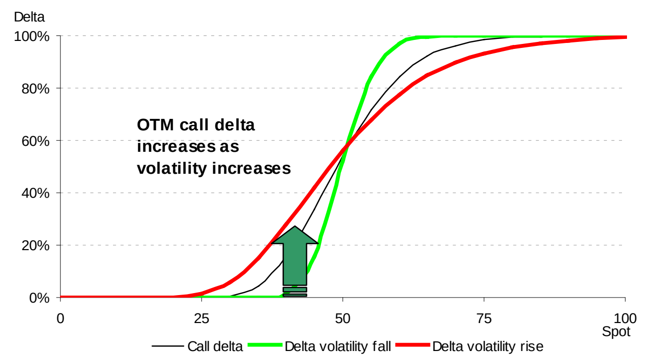

Source: Santander Investment Bolsa estimates.

MOVEMENT OF IMPLIED VOL SURFACE CHANGES DELTA OF OPTIONS

We shall assume we are in a sticky local vol (or jumpy vol) market, ie, volatility rises if markets fall, and a trader is trading skew using a long OTM put and short OTM call (ie, a risk reversal). As the delta of OTM options increases in value if implied volatility increases, and vice versa, the delta hedging of the long skew position is impacted by the movement in volatility surfaces.

Figure 126. Delta of Long Put if Market Falls (Local Vol)

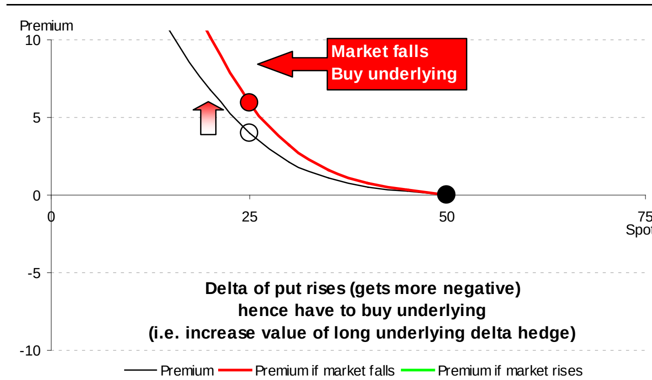

Delta of Short Call if Market Rises (Local Vol)

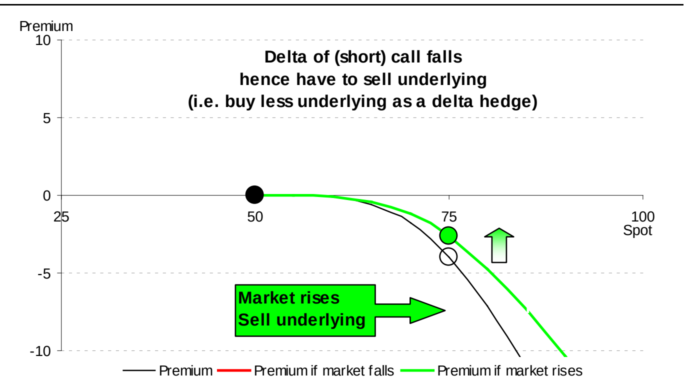

Source: Santander Investment Bolsa estimates.

Trader needs to buy stock (or futures) if market declines

If there is a decline in spot, the volatility of the long put (which is now more ATM and the primary driver of value) increases. This causes the delta of the position to decrease (absolute delta of put increases and, as delta of put is negative, the delta decreases). A trader has to buy more stock (or futures) than expected in order to compensate for this change, as shown on the left of Figure 126 above.

Conversely, trader needs to sell stock (or futures) if the market rises

The opposite trade occurs if markets rise as, for an increase in spot, the volatility of the short call (which is now more ATM and the primary driver of value) decreases. This causes the delta of the position to increase (delta of call decreases as delta of short call increases). Traders have to sell more stock (or futures) than they expect to compensate for this change (as shown on the right of Figure 126 above), which is the reverse trade of that which occurs for a decline in the market.

Figure 127. Delta Hedging Due to Trading Skew

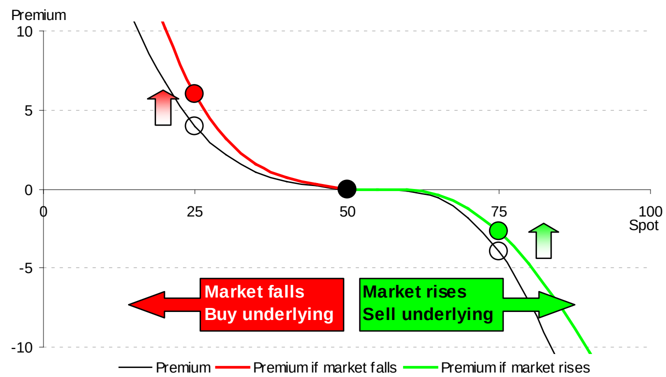

Delta Hedging Due to Trading Gamma

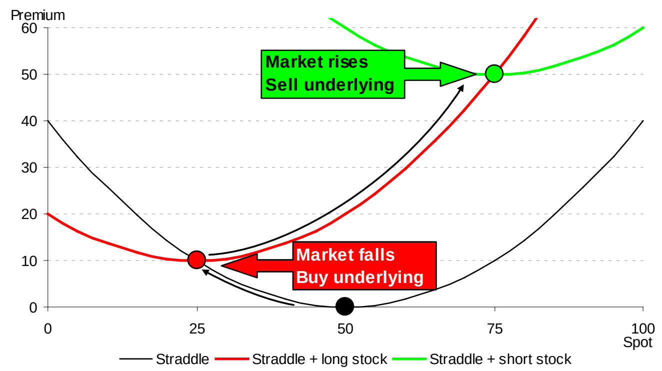

Source: Santander Investment Bolsa estimates.

DELTA HEDGING SKEW IS SIMILAR TO DELTA HEDGING GAMMA

Trading skew involves delta hedging the same way as trading gamma

Let us assume a negative correlation between spot and volatility (ie, for sticky local volatility or jumpy volatility) and that a trader is initially delta hedged[42] and intends to remain so. The movement of the volatility surface means the trader has to buy more stock (or futures) than he expects if markets fall and sell more stock (or futures) if markets rise. Buying low and selling high locks in the profit from the long skew position. This trade is identical to delta hedging a long gamma position, which can be seen in Figure 127 above.

If there is a positive relationship between spot and volatility (ie, a sticky delta volatility regime), then the reverse trade occurs with stock (or futures) being sold if markets decline and bought if markets rise. For sticky delta regimes, a long skew position is similar to being short gamma (and hence very unprofitable, given skew theta has to be paid as well).

[42] A long put and short call risk reversal would be delta hedged with a long stock (or futures) position.

MATHEMATICALLY, SKEW TRADING IS SIMILAR TO GAMMA TRADING

It is possible to show mathematically the relationship between skew trading and gamma trading if one assumes a correlation between spot and volatility. Vanna, the rate of change in vega for a change in spot (dVega/dSpot) measures the size of a skew position. This can be seen intuitively from the arguments above; as markets decline, the OTM put becomes more ATM and hence the primary driver of value. It is this change in vega (long put dominating the short call) for a change in spot, that causes volatility surface re-marks to be profitable for skew trading. Vanna is not only equal to dVega/dSpot, but is also equal to dDelta/dVol[43]. The equations below show that this relationship, when combined with spot being correlated to volatility, links skew and gamma trading.

Vanna = dDelta/dVol (and = dVega/dSpot)

As Vol α Spot

⇨ Vanna α dDelta/dSpot

As Gamma = dDelta/dSpot

⇨ Vanna α Gamma

Therefore, gamma can be considered to be second order gamma due to the negative correlation between volatility and spot.

[43] The proof of this relationship is outside of the scope of this publication

SKEW THETA PAYS FOR SKEW, GAMMA THETA PAYS FOR GAMMA

In order to break down an option's profit into volatility and skew, the total theta paid needs to be separated into gamma theta and skew theta. Gamma theta pays for gamma (or volatility) while skew theta pays for skew. We note that skew across the term structure can be compared with each other if weighted by the square root of time. As skew is measured by vanna, skew theta should therefore be compared to power vanna (vanna weighted by the square root of time) to identify skew trading opportunities. This is equivalent to comparing gamma to gamma theta. The method for calculating skew theta is given below.

Total theta = gamma theta + skew theta (all measured in theta per units of cash gamma)

Cash (or dollar) gamma = $\gamma \times S^2 / 100$ = notional cash value bought (or sold) per 1% spot move

GAMMA THETA IDENTICAL FOR ALL OPTIONS IF IMPLIEDS IDENTICAL

Gamma theta is the cost (or income) from a long (or short) gamma position. To calculate the cost of gamma, we shall assume an index has a volatility of 20% for all strikes and maturities. We shall ignore interest rates, dividends and borrowing costs and assume spot is currently at 3000pts.

Figure 128. Theta (per Year) of Index with 20% Implied

Strike	3 Months	1 Year	4 Years
80%	-0.070	-0.227	-0.178
90%	-0.517	-0.390	-0.213
100%	-0.949	-0.473	-0.233
110%	-0.632	-0.442	-0.237
120%	-0.197	-0.342	-0.230

Cash Gamma per 1% Move of Index with 20% Implied

Strike	3 Months	1 Year	4 Years
80%	9	29	22
90%	65	49	27
100%	120	60	29
110%	80	56	30
120%	25	43	29

Source: Santander Investment Bolsa estimates.

Both gamma and theta are high for short-dated ATM options

As can be seen in Figure 128, both cash gamma and theta are highest for near-dated and ATM options. When the cost per unit of cash gamma is calculated, it is identical for all strikes and expiries as the implied volatility is 20% for them all. We shall define the ATM theta cost per unit of cash gamma to be gamma theta (in units of 1 million cash gamma to have a reasonably sized number). This is, essentially, the values on the left in Figure 128 divided by the values on the right in Figure 128.

Figure 129. Theta per 1 million Cash Gamma (Gamma Theta) of Index with 20% Implied

Strike	3 Months	1 Year	4 Years
80%	-7,937	-7,937	-7,937
90%	-7,937	-7,937	-7,937
100% = Gamma Theta	-7,937	-7,937	-7,937
110%	-7,937	-7,937	-7,937
120%	-7,937	-7,937	-7,937

Source: Santander Investment Bolsa.

TERM STRUCTURE CHANGES GAMMA THETA BY MATURITY

In order to have a more realistic volatility surface we shall introduce positive sloping term structure, while keeping the implied volatility of one-year maturity options identical. As there is no skew in the surface, all the theta is solely due to the cost of gamma or gamma theta. The gamma theta is now lower for near-dated maturities, which is intuitively correct as near-dated implieds are now lower than the far-dated implieds.

Figure 130. Volatility Surface with Term Structure

Strike	3 Months	1 Year	4 Years
80%	19%	20%	21%
90%	19%	20%	21%
100%	19%	20%	21%
110%	19%	20%	21%
120%	19%	20%	21%

Theta per 1mn Cash Gamma (Gamma Theta)

Strike	3 Months	1 Year	4 Years
80%	-7,163	-7,937	-8,338
90%	-7,163	-7,937	-8,338
100%= Gamma Theta	-7,163	-7,937	-8,338
110%	-7,163	-7,937	-8,338
120%	-7,163	-7,937	-8,338

Source: Santander Investment Bolsa estimates.

SKEW MAKES IT MORE EXPENSIVE TO OWN PUTS THAN CALLS

If we introduce skew to the volatility surface we increase the cost of gamma for puts and decrease it for calls. This can be seen on the right of Figure 131; the ATM options have the same cost of gamma as before but the wings now have a different value.

Figure 131. Vol Surface with Skew and Term Structure

Strike	3 Months	1 Year	4 Years
80%	27%	24%	23%
90%	23%	22%	22%
100%	19%	20%	21%
110%	15%	18%	20%
120%	11%	16%	19%

Theta per 1mn Cash Gamma

Strike	3 Months	1 Year	4 Years
80%	-14,464	-11,429	-10,045
90%	-10,496	-9,603	-9,172
100%= Gamma Theta	-7,163	-7,937	-8,338
110%	-4,464	-6,429	-7,545
120%	-2,401	-5,079	-6,791

Source: Santander Investment Bolsa estimates.

SKEW THETA IS THE COST OF GOING LONG SKEW

As we have defined the theta paid for ATM option gamma (or gamma theta) as the fair price for gamma, the difference between this value and other options' cost of gamma is the cost of skew (or skew theta). Skew theta is therefore calculated by subtracting the cost of ATM gamma from all other options (and hence skew theta is zero for ATM options by definition).

Figure 132. Theta per 1mn Cash Gamma

Strike	3 Months	1 Year	4 Years
80%	-14,464	-11,429	-10,045
90%	**-10,496**	-9,603	-9,172
100%= Gamma Theta	**-7,163**	-7,937	-8,338
110%	-4,464	-6,429	-7,545
120%	-2,401	-5,079	-6,791

Skew Theta

Strike	3 Months	1 Year	4 Years
80%	-7,302	-3,492	-1,706
90%	**-3,333**	-1,667	-833
100%	0	0	0
110%	2,698	1,508	794
120%	4,762	2,857	1,548

Source: Santander Investment Bolsa estimates.

Example of skew theta calculation

The annual cost for a million units of cash gamma for three-month 90% strike options is €10,496, whereas ATM options only have to pay €7,163. The additional cost of being long 90% options (rather than ATM) is therefore €10,496 - €7,163 = €3,333. This additional €3,333 cost is the cost of being long skew, or skew theta.

Strikes lower than ATM suffer from skew theta

For low strike options there is a cost (negative sign) to owning the option and hence being long skew. High strike options benefit from an income of skew theta (which causes the lower cost of gamma) to compensate for being short skew (hence they have a positive sign).

VOLATILITY SLIDE THETA HAS A MINOR EFFECT ON SKEW TRADING

If one assumes volatility surfaces have relative time (one-year skew stays the same) rather than absolute time (ie, Dec14 skew stays the same) then one needs to take into account volatility slide theta (to factor in the increase in skew as the maturity of the option decreases). Volatility slide theta partly compensates for the cost of skew theta. For more details, see the section *Advanced (Practical or Shadow) Greeks* in the *Appendix*.

APPENDIX

LOCAL VOLATILITY

While Black-Scholes is the most popular method for pricing vanilla equity derivatives, exotic equity derivatives (and ITM American options) usually require a more sophisticated model. The most popular model after Black-Scholes is a local volatility model as it is the only completely consistent volatility model[44]. A local volatility model describes the instantaneous volatility of a stock, whereas Black-Scholes is the average of the instantaneous volatilities between spot and strike.

LOCAL VOLATILITY IS INSTANTANEOUS VOLATILITY OF UNDERLYING

Instantaneous volatility is the volatility of an underlying at any given local point, which we shall call the local volatility. We shall assume the local volatility is fixed and has a normal negative skew (higher volatility for lower spot prices). There are many paths from spot to strike and, depending on which path is taken, they will determine how volatile the underlying is during the life of the option (see Figure 133).

Figure 133. Different Paths between Spot and Strike

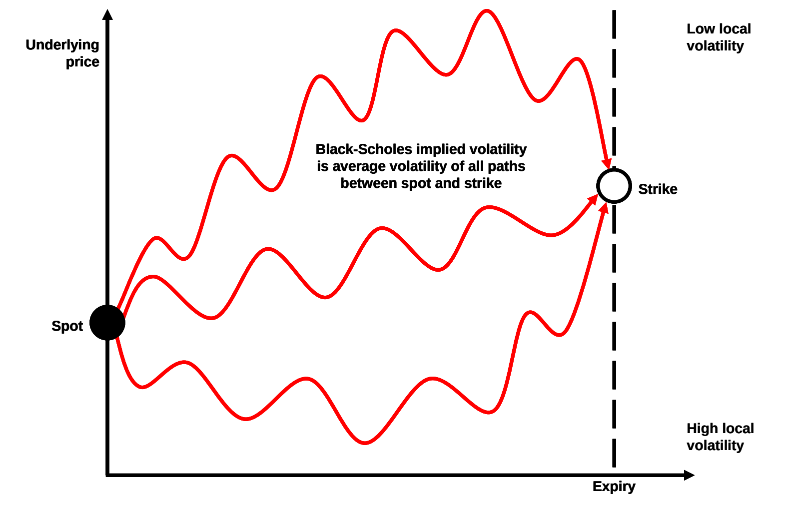

Source: Santander Investment Bolsa.

BLACK-SCHOLES VOLATILITY IS AVERAGE OF LOCAL VOLATILITIES

It is possible to calculate the local (or instantaneous) volatility surface from the Black-Scholes implied volatility surface. This is possible as the Black-Scholes implied volatility of an option is the average of all the paths between spot (ie, zero maturity ATM strike) and the maturity and strike of the option. A reasonable approximation is the average of all local volatilities on a direct straight-line path between spot and strike. For a normal relatively flat skew, this is simply the average of two values, the ATM local volatility and the strike local volatility.

---

[44] Strictly speaking, this is true only for deterministic models. However, as the expected volatility of non-deterministic models has to give identical results to a local volatility model to be completely consistent, they can be considered to be a 'noisy' version of a local volatility model.

Black-Scholes skew is half local volatility skew as it is the average

ATM implieds are identical for local vol and Black-Scholes, but local vol skew is twice Black-Scholes

If the local volatility surface has a 22% implied at the 90% strike, and 20% implied at the ATM strike, then the Black-Scholes implied volatility for the 90% strike is 21% (average of 22% and 20%). As ATM implieds are identical for both local and Black-Scholes implied volatility, this means that 90%-100% skew is 2% for local volatility but 1% for Black-Scholes. Local volatility skew is therefore twice the Black-Scholes skew.

ATM volatility is the same for both Black-Scholes and local volatility

For ATM implieds, the local volatility at the strike is equal to ATM, hence the average of the two identical numbers is simply equal to the ATM implied. For this reason, Black-Scholes ATM implied is equal to local volatility ATM implied.

LOCAL VOL IS THE ONLY COMPLETE CONSISTENT VOL MODEL

A local volatility model is complete (it allows hedging based only on the underlying asset) and consistent (does not contain a contradiction). It is often used to calculate exotic option implied volatilities to ensure the prices for these exotics are consistent with the values of observed vanilla options and hence prevent arbitrage. A local volatility model is the only complete consistent volatility model; a constant Black-Scholes volatility model (constant implied volatility for all strikes and expiries) can be considered to be a special case of a static local volatility model (where the local volatilities are fixed and constant for all strikes and expiries).

MEASURING HISTORICAL VOLATILITY

The implied volatility for a certain strike and expiry has a fixed value. There is, however, no single calculation for historical volatility. The number of historical days for the historical volatility calculation changes the calculation, in addition to the estimate of the drift (or average amount stocks are assumed to rise). There should, however, be no difference between the average daily or weekly historical volatility. We also examine different methods of historical volatility calculation, including close-to-close volatility and exponentially weighted volatility, in addition to advanced volatility measures such as Parkinson, Garman-Klass (including Yang-Zhang extension), Rogers and Satchell and Yang-Zhang. We also show that it is best to assume a zero drift assumption for close-to-close volatility, and that under this condition variance is additive.

DEFINITION OF VOLATILITY

Assuming that the probability distribution of the log returns of a particular security is normally distributed (or follows a normal 'bell-shape distribution'), volatility σ of that security can be defined as the standard deviation of the normal distribution of the log returns. As the mean absolute deviation is $\sqrt{(2/\pi)}$ (≈0.8) × volatility, the volatility can be thought of as c1.25× the expected percentage change (positive or negative) of the security.

$$\sigma = \text{standard deviation of log returns} \times \sqrt{1/\Delta t}$$

CLOSE-TO-CLOSE HISTORICAL VOLATILITY IS THE MOST COMMON

Volatility is defined as the annualised standard deviation of log returns. For historical volatility the usual measure is close-to-close volatility, which is shown below.

Log return = $x_i = \text{Ln}\left(\frac{c_i + d_i}{c_{i-1}}\right)$ where d_i = ordinary dividend and c_i is close price

Volatility[45] (not annualised) = $\sigma_x = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^2}$

where $\bar{x}$ = drift = Average (x_i)

Historical volatility calculation is an estimate from a sample

Historical volatility is calculated as the standard deviation of the log returns of a particular securities' time series. If the log returns calculation is based on daily data, we have to multiply this number by the square root of 252 (the number of trading days in a calendar year) in order to annualise the volatility calculation (as $\Delta t = 1/252$ hence $\sqrt{1/\Delta t} = \sqrt{252}$). As a general rule, to annualise the volatility calculation, regardless of the periodicity of the data, the standard deviation has to be multiplied by the square root of the number of days/weeks/months within a year (ie, $\sqrt{252}, \sqrt{52}, \sqrt{12}$).

$$\sigma_{Annualised} = \sigma_x \times \sqrt{\text{values in year}}$$

[45] We take the definition of volatility of John Hull in *Options, Futures and Other Derivatives* in which n day volatility uses n returns and n+1 prices. We note Bloomberg uses n prices and n-1 returns.

VARIANCE IS ADDITIVE IF ZERO MEAN IS ASSUMED

Frequency of returns in a year = F (eg, 252 for daily returns)

$$\sigma_{\text{Annualised}} = \sqrt{F} \times \sigma_x = \sqrt{F}\sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^2}$$

As $\bar{x} \approx 0$ if we assume zero average returns

$$\sigma_{\text{Annualised}} = \sqrt{F}\sqrt{\frac{1}{N}\sum_{i=1}^{N}x_i^{\,2}}$$

$$\sigma_{\text{Annualised}}^{\,2} = \frac{F}{N}\sum_{i=1}^{N}x_i^{\,2}$$

Now if we assume that the total sample N can be divided up into period 1 and period 2 where period 1 is the first M returns then:

$$\sigma_{\text{Total}}^{\,2} = \frac{F}{N_{\text{Total}}} \times \sum_{i=1}^{N}x_i^{\,2}$$

$$\sigma_{\text{Period 1}}^{\,2} = \frac{F}{N_{\text{Period 1}}} \times \sum_{i=1}^{M}x_i^{\,2} \text{ (where } N_{\text{Period 1}} = M\text{)}$$

$$\sigma_{\text{Period 2}}^{\,2} = \frac{F}{N_{\text{Period 2}}} \times \sum_{i=M+1}^{N}x_i^{\,2} \text{ (where } N_{\text{Period 2}} = N - M\text{)}$$

then

$$\sigma_{\text{Total}}^{\,2} = \frac{F}{N_{\text{Total}}} \times \sum_{i=1}^{N}x_i^{\,2} = \frac{F}{N_{\text{Total}}}\left(\sum_{i=1}^{M}x_i^{\,2} + \sum_{i=M+1}^{N}x_i^{\,2}\right)$$

$$\sigma_{\text{Total}}^{\,2} = \frac{F}{N_{\text{Total}}} \times \sum_{i=1}^{M}x_i^{\,2} + \frac{F}{N_{\text{Total}}} \times \sum_{i=M+1}^{N}x_i^{\,2}$$

$$\sigma_{\text{Total}}^{\,2} = \frac{N_{\text{Period 1}}}{N_{\text{Total}}}\left(\frac{F}{N_{\text{Period 1}}} \times \sum_{i=1}^{M}x_i^{\,2}\right) + \frac{N_{\text{Period 2}}}{N_{\text{Total}}}\left(\frac{F}{N_{\text{Period 2}}} \times \sum_{i=M+1}^{N}x_i^{\,2}\right)$$

$$\sigma_{\text{Total}}^{\,2} = \frac{N_{\text{Period 1}}}{N_{\text{Total}}}\sigma_{\text{Period 1}}^{\,2} + \frac{N_{\text{Period 2}}}{N_{\text{Total}}}\sigma_{\text{Period 2}}^{\,2}$$

Hence variance is additive (when weighted by the time in each period / total time)

BEST TO ASSUME ZERO DRIFT FOR VOLATILITY CALCULATION

For relatively short time periods (daily, weekly), the drift should be close to zero and can be ignored

The calculation for standard deviation calculates the deviation from the average log return (or drift). This average log return has to be estimated from the sample, which can cause problems if the return over the period sampled is very high or negative. As over the long term very high or negative returns are not realistic, the calculation of volatility can be corrupted by using the sample log return as the expected future return. For example, if an underlying rises 10% a day for ten days, the volatility of the stock is zero (as there is zero deviation from the 10% average return). This is why volatility calculations are normally more reliable if a zero return is assumed. In theory, the expected average value of an underlying at a future date should be the value of the forward at that date. As for all normal interest rates (and dividends, borrow cost) the forward return should be close to 100% (for any reasonable sampling frequency, ie, daily/weekly/monthly). Hence, for simplicity reasons it is easier to assume a zero log return as Ln(100%) = 0.

WHICH HISTORICAL VOLATILITY SHOULD I USE?

When examining how attractive the implied volatility of an option is, investors will often compare it to historical volatility. However, historical volatility needs two parameters.

- Length of time (eg, number of days/weeks/months)
- Frequency of measurement (eg, daily/weekly)

LENGTH OF TIME FOR HISTORICAL VOLATILITY

Historical volatility should be a multiple of 3 months to have a constant number of quarterly reporting periods

Choosing the historical volatility number of days is not a trivial choice. Some investors believe the best number of days of historical volatility to look at is the same as the implied volatility of interest. For example, one-month implied should be compared to 21 trading day historical volatility (and three-month implied should be compared to 63-day historical volatility, etc). While an identical duration historical volatility is useful to arrive at a realistic minimum and maximum value over a long period of time, it is not always the best period of time to determine the fair level of long-dated implieds. This is because volatility mean reverts over a period of c8 months. Using historical volatility for periods longer than c8 months is not likely to be the best estimate of future volatility (as it could include volatility caused by earlier events, whose effect on the market has passed). Arguably a multiple of three months should be used to ensure that there is always the same number of quarterly reporting dates in the historical volatility measure. Additionally, if there has been a recent jump in the share price that is not expected to reoccur, the period of time chosen should try to exclude that jump.

The best historical volatility period does not have to be the most recent

If there has been a rare event which caused a volatility spike, the best estimate of future volatility is not necessary the current historical volatility. A better estimate could be the past historical volatility when an event that caused a similar volatility spike occurred. For example, the volatility post credit crunch could be compared to the volatility spike after the Great Depression or during the bursting of the tech bubble.

FREQUENCY OF HISTORICAL VOLATILITY

While historical volatility can be measured monthly, quarterly or yearly, it is usually measured daily or weekly. Normally, daily volatility is preferable to weekly volatility as five times as many data points are available. However, if volatility over a long period of time is being examined between two different markets, weekly volatility could be the best measure to reduce the influence of different public holidays (and trading hours[46]). If stock price returns are independent, then the daily and weekly historical volatility should on average be the same. If stock price returns are not independent, there could be a difference. Autocorrelation is the correlation between two different returns so independent returns have an autocorrelation of 0%.

Trending markets imply weekly volatility is greater than daily volatility

With 100% autocorrelation, returns are perfectly correlated (ie, trending markets). Should autocorrelation be -100% correlated, then a positive return is followed by a negative return (mean reverting or range trading markets). If we assume markets are 100% daily correlated with a 1% daily return, this means the weekly return is 5%. The daily volatility is therefore c16% (1% × √252), while the weekly volatility of c35% (5% × √52) is more than twice as large.

Figure 134. Stock Price with 100% Daily Autocorrelation **Stock Price with -100% Daily Autocorrelation**

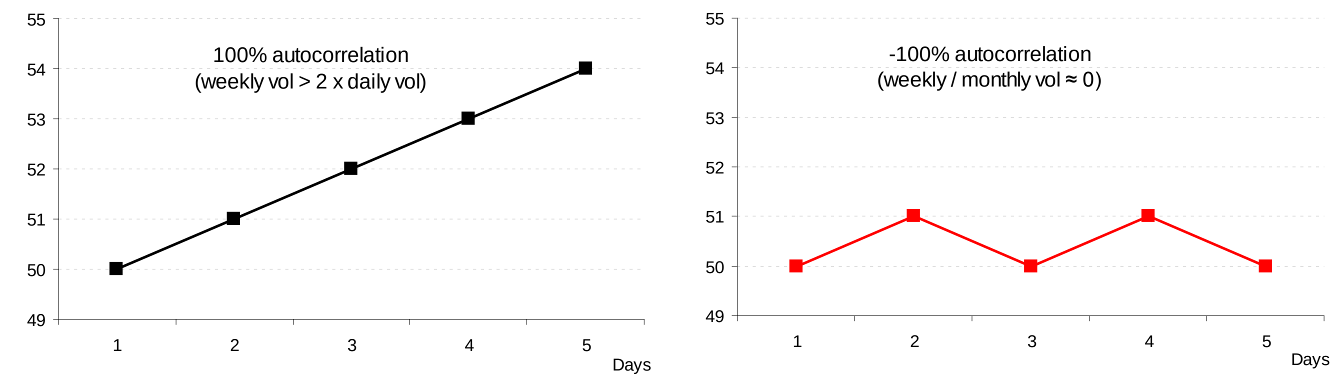

Source: Santander Investment Bolsa estimates.

High market share of high frequency trading should prevent autocorrelation

Historically (decades ago), there could have been positive autocorrelation due to momentum buying, but once this became understood this effect is likely to have faded. Given the current high market share of high frequency trading (accounting for up to three-quarters of US equity trading volume), it appears unlikely that a simple trading strategy such as ‘buy if security goes up, sell if it goes down’ will provide above-average returns over a significant period of time[47].

Panicked markets could cause temporary negative autocorrelation

While positive autocorrelation is likely to be arbitraged out of the market, there is evidence that markets can overreact at times of stress as market panic (rare statistical events can occur under the weak form of efficient market hypotheses). During these events human traders and some automated trading systems are likely to stop trading (as the event is rare, the correct response is unknown), or potentially exaggerate the trend (as positions get ‘stopped out’ or to follow the momentum of the move). A strategy that is long daily variance and short weekly variance will therefore usually give relatively flat returns, but occasionally give a positive return.

[46] Advanced volatility measures could be used to remove part of the effect of different trading hours.
[47] Assuming there are no short selling restrictions.

INTRADAY VOLATILITY IS NOT CONSTANT

Advanced volatility measures should be used by traders wishing to take into account intraday prices

For most markets, intraday volatility is greatest just after the open (as results are often announced around the open) and just before the close (performance is often based upon closing prices). Intraday volatility tends to sag in the middle of the day due to the combination of a lack of announcements and reduced volumes/liquidity owing to lunch breaks. For this reason, using an estimate of volatility more frequent than daily tends to be very noisy. Traders who wish to take into account intraday prices should instead use an advanced volatility measure.

Figure 135. Intraday Volatility

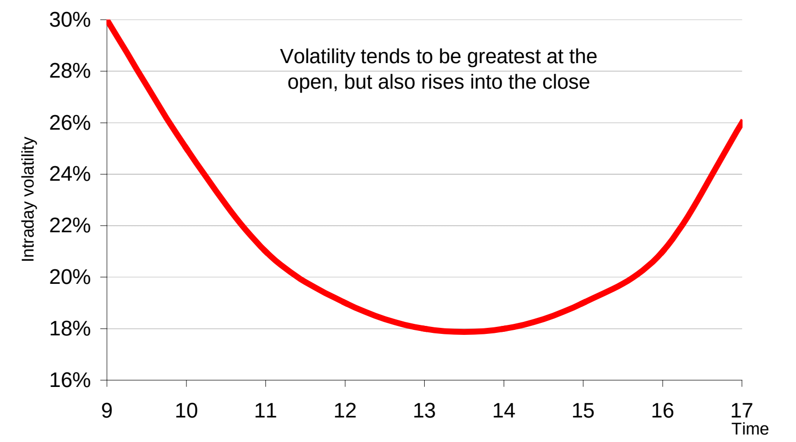

Source: Santander Investment Bolsa.

EXPONENTIALLY WEIGHTED VOLATILITIES ARE RARELY USED

Exponentially weighted moving average can be used to reduce effect of spikes in volatility disappearing

An alternate measure could be to use an exponentially weighted moving average model, which is shown below. The parameter λ is between zero (effectively one-day volatility) and one (ignore current vol and keep vol constant). Normally, values of c0.9 are used. Exponentially weighted volatilities are rarely used, partly due to the fact they do not handle regular volatility-driving events such as earnings very well. Previous earnings jumps will have least weight just before an earnings date (when future volatility is most likely to be high) and most weight just after earnings (when future volatility is most likely to be low). It could, however, be of some use for indices.

$$\sigma_i^2 = \lambda \sigma_{i-1}^2 + (1-\lambda) x_i^2$$

Exponentially weighted volatility avoids volatility collapse of historic volatility

Exponential volatility has the advantage over standard historical volatility in that the effect of a spike in volatility gradually fades (as opposed to suddenly disappearing causing a collapse in historic volatility). For example, if we are looking at the historical volatility over the past month and a spike in realised volatility suddenly occurs the historical volatility will be high for a month, then collapse. Exponentially weighted volatility will rise at the same time as historical volatility and then gradually decline to lower levels (arguably in a similar way to how implied volatility spikes, then mean reverts).

ADVANCED VOLATILITY MEASURES

Volatility measures can use open, high and low prices in addition to closing price

Close-to-close volatility is usually used as it has the benefit of using the closing auction prices only. Should other prices be used, then they could be vulnerable to manipulation or a 'fat fingered' trade. However, a large number of samples need to be used to get a good estimate of historical volatility, and using a large number of closing values can obscure short-term changes in volatility. There are, however, different methods of calculating volatility using some or all of the open (O), high (H), low (L) and close (C). The methods are listed in order of their maximum efficiency (close-to-close variance divided by alternative measure variance).

- **Close to close (C).** The most common type of calculation that benefits from only using reliable prices from closing auctions. By definition its efficiency is one at all times.

- **Parkinson (HL).** As this estimate only uses the high and low price for an underlying, it is less sensitive to differences in trading hours. For example, as the time of the EU and US closes are approximately half a trading day apart, they can give very different returns. Using the high and low means the trading over the whole day is examined, and the days overlap. As it does not handle jumps, on average it underestimates the volatility, as it does not take into account highs and lows when trading does not occur (weekends, between close and open). Although it does not handle drift, this is usually small. The Parkinson estimate is up to 5.2 times more efficient than the close-to-close estimate. While other measures are more efficient based on simulated data, some studies have shown it to be the best measure for actual empirical data.

- **Garman-Klass (OHLC).** This estimate is the most powerful for stocks with Brownian motion, zero drift and no opening jumps (ie, opening price is equal to closing price of previous period). While it is up to 7.4 times as efficient as the close to close estimate, it also underestimates the volatility (as like Parkinson it assumes no jumps).

- **Rogers-Satchell (OHLC).** The efficiency of the Rogers-Satchell estimate is similar to that for Garman-Klass; however, it benefits from being able to handle non-zero drift. Opening jumps are not handled well though, which means it underestimates the volatility.

- **Garman-Klass Yang-Zhang extension (OHLC).** Yang-Zhang extended the Garman-Klass method that allows for opening jumps hence it is a fair estimate, but does assume zero drift. It has an efficiency of eight times the close-to-close estimate.

- **Yang-Zhang (OHLC).** The most powerful volatility estimator which has minimum estimation error. It is a weighted average of Rogers-Satchell, the close-open volatility and the open-close volatility. It is up to a maximum of 14 times as efficient (for two days of data) as the close-to-close estimate.

Figure 136. Summary of Advanced Volatility Estimates

Estimate	Prices Taken	Handle Drift?	Handle Overnight Jumps?	Efficiency (max)
Close to close	C	No	No	**1**
Parkinson	HL	No	No	**5.2**
Garman-Klass	OHLC	No	No	**7.4**
Rogers-Satchell	OHLC	Yes	No	**8**
Garman-Klass Yang-Zhang ext.	OHLC	No	Yes	**8**
Yang-Zhang	OHLC	Yes	Yes	**14**

Source: Santander Investment Bolsa.

EFFICIENCY AND BIAS DETERMINE BEST VOLATILITY MEASURE

Standard close-to-close is best for large samples, Yang-Zhang is best for small samples

There are two measures that can be used to determine the quality of a volatility measure: efficiency and bias. Generally, for small sample sizes the Yang-Zhang measure is best overall, and for large sample sizes the standard close to close measure is best.

- **Efficiency.** Efficiency$(\sigma_x^2) = \frac{\sigma_{cc}^2}{\sigma_x^2}$ where σ_x is the volatility of the estimate and σ_{cc} is the volatility of the standard close to close estimate.
- **Bias.** Difference between the estimated variance and the average (ie, integrated) volatility.

Efficiency measures the volatility of the estimate

The efficiency describes the variance, or volatility of the estimate. The efficiency is dependent on the number of samples, with efficiency decreasing the more samples there are (as close-to-close will converge and become less volatile with more samples). The efficiency is the theoretical maximum performance against an idealised distribution, and with real empirical data a far smaller benefit is usually seen (especially for long time series). For example, while the Yang-Zhang based estimators deal with overnight jumps if the jumps are large compared to the daily volatility the estimate will converge with the close-to-close volatility and have an efficiency close to one.

Close-to-close volatility should use at least five samples (and ideally 20 or more)

The variance of the close-to-close volatility can be estimated as a percentage of the actual variance by the formula 1/(2N) where N is the number of samples. This is shown in Figure 137 below and demonstrates that at least five samples are needed (or the estimate has a variance of over 10%) and that only marginal extra accuracy is gained for each additional sample above 20.

Figure 137. Variance of Close-To-Close Volatility/Actual Variance

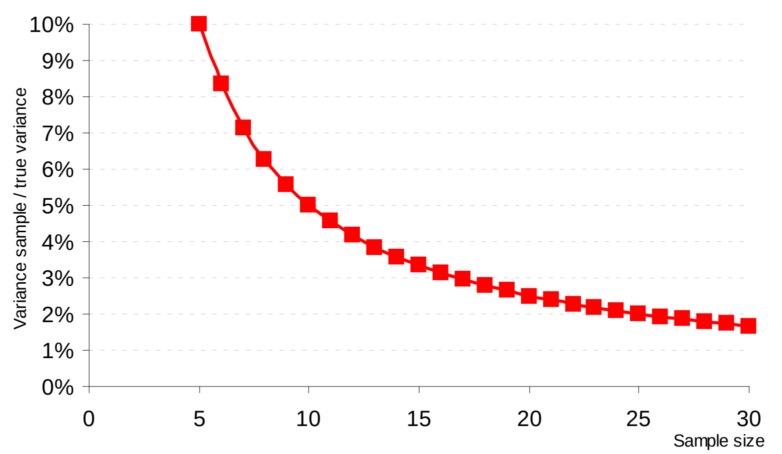

Source: Santander Investment Bolsa.

Bias depends on the type of distribution of the underlying

Bias can be positive or negative depending on the distribution

While efficiency (how volatile the measure is) is important, so too is bias (whether the measure is, on average, too high or low). Bias depends on the sample size, and the type of distribution the underlying security has. Generally, the close-to-close volatility estimator is too big[48] (as it does not model overnight jumps), while alternative estimators are too small (as they assume continuous trading, and discrete trading will have a smaller difference between the maximum and minimum). The key variables that determine the bias are:

- **Sample size.** As the standard close-to-close volatility measure suffers with small sample sizes, this is where alternative measures perform best (the highest efficiency is reached for only two days of data).

- **Volatility of volatility.** While the close-to-close volatility estimate is relatively insensitive to a changing volatility (vol of vol), the alternative estimates are far more sensitive. This bias increases the more vol of vol increases (ie, more vol of vol means a greater underestimate of volatility).

Approximately 1/6 of total volatility occurs overnight

- **Overnight jumps between close and open.** Approximately one-sixth of equity volatility occurs outside the trading day (and approximately twice that amount for ADRs). Overnight jumps cause the standard close-to-close estimate to overestimate the volatility, as jumps are not modelled. Alternative estimates that do not model jumps (Parkinson, Garman Klass and Rogers-Satchell) underestimate the volatility. Yang-Zhang estimates (both Yang-Zhang extension of Garman Klass and the Yang-Zhang measure itself) will converge with standard close-to-close volatility if the jumps are large compared to the overnight volatility.

- **Drift of underlying.** If the drift of the underlying is ignored as it is for Parkinson and Garman Klass (and the Yang Zhang extension of Garman Glass), then the measure will overestimate the volatility. This effect is small for any reasonable drifts (ie, if we are looking at daily, weekly or monthly data).

- **Correlation daily volatility and overnight volatility.** While Yang-Zhang measures deal with overnight volatility, there is the assumption that overnight volatility and daily volatility are uncorrelated. Yang-Zhang measures will underestimate volatility when there is a correlation between daily return and overnight return (and vice versa), but this effect is small.

Variance, volatility and gamma swaps should look at standard volatility (or variance)

As the payout of variance, volatility and gamma swaps are based on close-to-close prices, the standard close-to-close volatility (or variance) should be used for comparing their price against realised. Additionally, if a trader only hedges at the close (potentially for liquidity reasons) then again the standard close-to-close volatility measure should be used.

[48] Compared to integrated volatility.

CLOSE-TO-CLOSE

As the average is taken from the sample, close-to-close volatility has N-1 degrees of freedom

The simplest volatility measure is the standard close-to-close volatility. We note that the volatility should be the standard deviation multiplied by √N/(N-1) to take into account the fact we are sampling the population (or take standard deviation of the sample)[49]. We ignored this in the earlier definition as for reasonably large n it √N/(N-1) is roughly equal to one.

Standard dev of x = $s_x = \sqrt{\frac{F}{N}}\sqrt{\sum_{i=1}^{N}(x_i - \bar{x})^2}$

As $\sigma = \sqrt{\sigma^2} = \sqrt{E(s^2)} < E(\sqrt{s^2}) = E(s)$ by Jensens´s inequality

Volatility = $\sigma_x = s_x \times \sqrt{\frac{N}{N-1}}$

Volatility$_{\text{close to close}} = \sigma_{cc} = \sqrt{\frac{F}{N-1}}\sqrt{\sum_{i=1}^{N}(x_i - \bar{x})^2} = \sqrt{\frac{F}{N-1}}\sqrt{\sum_{i=1}^{N}\left(Ln(\frac{c_i}{c_{i-1}})\right)^2}$ with zero drift

PARKINSON

The first advanced volatility estimator was created by Parkinson in 1980, and instead of using closing prices it uses the high and low price. One drawback of this estimator is that it assumes continuous trading, hence it underestimates the volatility as potential movements when the market is shut are ignored.

Volatility$_{\text{Parkinson}} = \sigma_P = \sqrt{\frac{F}{N}}\sqrt{\frac{1}{4\,Ln(2)}\sum_{i=1}^{N}\left(Ln(\frac{h_i}{l_i})\right)^2}$

GARMAN-KLASS

Later in 1980 the Garman-Klass volatility estimator was created. It is an extension of Parkinson which includes opening and closing prices (if opening prices are not available the close from the previous day can be used instead). As overnight jumps are ignored the measure underestimates the volatility.

Volatility$_{\text{Garman-Klass}} = \sigma_{GK} = \sqrt{\frac{F}{N}}\sqrt{\sum_{i=1}^{N}\frac{1}{2}\left(Ln(\frac{h_i}{l_i})\right)^2 - (2Ln(2)-1)\left(Ln(\frac{c_i}{o_i})\right)^2}$

[49] As the formula for standard deviation has N-1 degrees of freedom (as we subtract the sample average from each value of x)

ROGERS-SATCHELL

All of the previous advanced volatility measures assume the average return (or drift) is zero. Securities that have a drift, or non-zero mean, require a more sophisticated measure of volatility. The Rogers-Satchell volatility created in the early 1990s is able to properly measure the volatility for securities with non-zero mean. It does not, however, handle jumps; hence, it underestimates the volatility.

$$\text{Volatility}_{\text{Rogers-Satchell}} = \sigma_{RS} = \sqrt{\frac{F}{N}}\sqrt{\sum_{i=1}^{N} Ln(\frac{h_i}{c_i})Ln(\frac{h_i}{o_i}) + Ln(\frac{l_i}{c_i})Ln(\frac{l_i}{o_i})}$$

GARMAN-KLASS YANG-ZHANG EXTENSION

Yang-Zhang modified the Garman-Klass volatility measure in order to let it handle jumps. The measure does assume a zero drift; hence, it will overestimate the volatility if a security has a non-zero mean return. As the effect of drift is small, the fact continuous prices are not available usually means it underestimates the volatility (but by a smaller amount than the previous alternative measures).

$$\text{Volatility}_{\text{GKYZ}} = \sigma_{\text{GKYZ}} = \sqrt{\frac{F}{N}}\sqrt{\sum_{i=1}^{N}\left(Ln(\frac{o_i}{c_{i-1}})\right)^2 + \frac{1}{2}\left(Ln(\frac{h_i}{l_i})\right)^2 - (2Ln(2)-1)\left(Ln(\frac{c_i}{o_i})\right)^2}$$

YANG-ZHANG

Yang-Zhang is the sum of overnight volatility, and a weighted average of Rogers-Satchell and open-to-close volatility

In 2000 Yang-Zhang created a volatility measure that handles both opening jumps and drift. It is the sum of the overnight volatility (close-to-open volatility) and a weighted average of the Rogers-Satchell volatility and the open-to-close volatility. The assumption of continuous prices does mean the measure tends to slightly underestimate the volatility.

$$\text{Volatility}_{\text{Yang-Zhang}} = \sigma_{YZ} = \sqrt{\frac{F}{N}}\sqrt{\sigma_{\text{overnight volatility}}{}^2 + k\,\sigma_{\text{open to close volatility}}{}^2 + (1-k)\sigma_{RS}{}^2}$$

where $k = \dfrac{0.34}{1.34 + \dfrac{N+1}{N-1}}$

$$\sigma_{\text{overnight volatility}}{}^2 = \frac{1}{N-1}\sum_{i=1}^{N}\left[Ln(\frac{o_i}{c_{i-1}}) - \overline{Ln(\frac{o_i}{c_{i-1}})}\right]^2$$

$$\sigma_{\text{open to close volatility}}{}^2 = \frac{1}{N-1}\sum_{i=1}^{N}\left[Ln(\frac{c_i}{o_i}) - \overline{Ln(\frac{c_i}{o_i})}\right]^2$$

PROOF VARIANCE SWAPS CAN BE HEDGED BY LOG CONTRACT (=1/K^2)

A log contract is a portfolio of options of all strikes (K) weighted by 1/K^2. When this portfolio of options is delta hedged on the close, the payoff is identical to the payoff of a variance swap. We prove this relationship and hence show that the volatility of a variance swap can be hedged with a static position in a log contract.

PORTFOLIO OF OPTIONS WITH CONSTANT VEGA WEIGHTED 1/K^2

In order to prove that a portfolio of options with flat vega has to be weighted 1/K^2, we will define the variable x to be K/S (strike K divided by spot S). With this definition and assuming zero interest rates, the standard Black-Scholes formula for vega of an option simplifies to:

Vega of option = $\tau \times S \times f(x, v)$

where

x = K / S (strike a ratio of spot)

τ = time to maturity

v = $\sigma^2 \tau$ (total variance)

$$f(x, v) = \frac{1}{\sqrt{2\pi}} \times e^{\frac{-d_1^2}{2}}$$

$$d_1 = \frac{\mathrm{Ln}(\frac{1}{x}) + \frac{v}{2}}{\sqrt{v}}$$

If we have a portfolio of options where the weight of each option is w(K), then the vega of the portfolio of options V(S) is:

$$V(S) = \tau \int_{K=0}^{\infty} w(K) \times S \times f(x, v) dK$$

As K = xS this means dK / dx = S, hence dK = S × dx and we can change variable K for x.

$$V(S) = \tau \int_{x=0}^{\infty} w(xS) \times S^2 \times f(x, v) dx$$

In order for the portfolio of options to have a constant vega – no matter what the level of spot – dV(S)/dS has to be equal to zero.

$$\frac{dV}{dS} = \tau \int_{x=0}^{\infty} \frac{d}{dS}\left[S^2 w(xS)\right] \times f(x,v)dx = 0$$

And by the chain rule:

$$\Rightarrow \quad \tau \int_{x=0}^{\infty} \left[2Sw(xS) + S^2 \frac{d}{dS} w(xS) \right] \times f(x,v)dx = 0$$

$$\Rightarrow \quad \tau \int_{x=0}^{\infty} S\left[2w(xS) + S \frac{d}{dS} w(xS) \right] \times f(x,v)dx = 0$$

As d/dS = (d/dK) × (dK/dS), and dK/dS = x

$$\Rightarrow \quad \tau \int_{x=0}^{\infty} S\left[2w(xS) + xS \frac{d}{dK} w(xS) \right] \times f(x,v)dx = 0$$

As K = xS

$$\Rightarrow \quad \tau \int_{x=0}^{\infty} S\left[2w(xS) + K \frac{d}{dK} w(K) \right] \times f(x,v)dx = 0$$

$$\Rightarrow \quad 2w + K \frac{d}{dK} w(K) = 0 \text{ for all values of S}$$

$$\Rightarrow \quad w(K) = \frac{\text{constant}}{K^2}$$

PROOF VARIANCE SWAP NOTIONAL = VEGA/2σ

For small differences between the future volatility and current (implied) swap volatility, the payout of a volatility swap can be approximated by a variance swap. We show how the difference in their notionals should be weighted by 2σ.

Proof that Variance Swap Notional = vega/2σ

We intend to calculate the relative size (Z) of the variance swap notional compared to volatility swap notional (volatility swap notional = vega by definition) so they have a similar payout (for small differences between realised and implied volatility).

$$\text{Notional}_{\text{variance swap}} \approx Z \times \text{Notional}_{\text{volatility swap}}$$

⇨ $(\sigma_F - \sigma_S) \approx Z\,(\sigma_F^2 - \sigma_S^2)$

where:

σ_F = future volatility (that occurs over the life of contract)

σ_S = swap rate volatility (fixed at the start of contract)

As there is a small difference between future (realised) volatility and swap rate (implied) volatility, then we can define $\sigma_F = \sigma_S + x$ where x is small.

⇨ $((\sigma_S + x) - \sigma_S) \approx Z\,((\sigma_S + x)^2 - \sigma_S^2)$ for simplification we shall replace σ_S with σ

⇨ $x \approx Z\,((\sigma^2 + 2\sigma x + x^2) - \sigma^2)$

⇨ $x \approx Z\,(2\sigma x + x^2)$

⇨ $1 \approx Z\,(2\sigma + x)$ and as x is small

⇨ $1/2\sigma \approx Z$

Hence $\text{Notional}_{\text{variance swap}}$ = vega / 2σ (as vega = $\text{Notional}_{\text{volatility swap}}$)

MODELLING VOLATILITY SURFACES

There are a variety of constraints on the edges of a volatility surface, and this section details some of the most important constraints from both a practical and theoretical point of view. We examine the considerations for very short-dated options (a few days or weeks), options at the wings of a volatility surface and very long-dated options.

IMPLIED VOLATILITY IS LESS USEFUL FOR NEAR-DATED OPTIONS

Options that only have a few days or a few weeks to expiry have a very small premium. For these low-value options, a relatively small change in price will equate to a relatively large change in implied volatility. This means the implied volatility bid-offer arbitrage channel is wider, and hence less useful. The bid-offer spread is more stable in cash terms for options of different maturity, so shorter-dated options should be priced more by premium rather than implied volatility.

Need to price short-dated options with a premium after a large collapse in the market

Near-dated options need to take premium, as well as implied volatility, into account

If there has been a recent dip in the market, there is a higher than average probability that the markets could bounce back to their earlier levels. The offer of short-dated ATM options should not be priced at a lower level than the size of the decline. For example, if markets have dropped 5%, then a one-week ATM call option should not be offered for less than c5% due to the risk of a bounce-back.

SKEW SHOULD DECAY BY SQUARE ROOT OF TIME

The payout of a put spread (and call spread) is always positive; hence, it should always have a positive cost. If it was possible to enter into a long put (or call) spread position for no cost (or potentially earning a small premium), any rational investor would go long as large a position as possible and earn risk-free profits (as the position cannot suffer a loss). A put spread will have a negative cost if the premium earned by selling the lower strike put is more than the premium of the higher strike put bought. This condition puts a cap on how negative skew can be: for high (negative) skew, the implied of the low strike put could be so large the premium is too high (ie, more than the premium of higher strike puts). The same logic applies for call spreads, except this puts a cap on positive skew (ie, floor on negative skew). As skew is normally negative, the condition on put spreads (see figure below on the left) is usually the most important. As time increases, it can be shown that the cap and floor for skew (defined as the gradient of first derivative of volatility with respect to strike, which is proportional to 90%-100% skew) decays by roughly the square root of time. This gives a mathematical basis for the 'square root of time rule' used by traders.

Figure 138. Put Spread

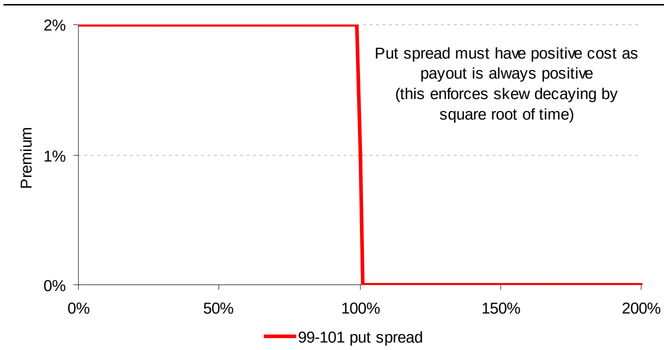

Ratio put spread

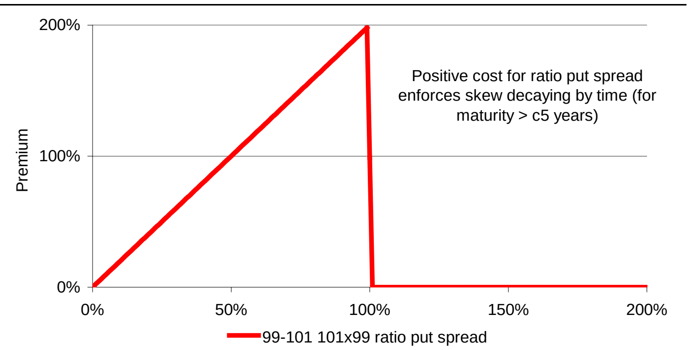

Source: Santander Investment Bolsa estimates.

Far-dated skew should decay by time for long maturities (c5 years)

In theory skew should decay by time (not square root of time)

It is possible to arrive at a stronger limit to the decay of skew by considering leveraged ratio put spreads (see chart above on the right). For any two strikes A and B (assume A<B), then the payout of going long A× puts with strike B, and going short B× puts with strike A creates a ratio put spread whose value cannot be less than zero. This is because the maximum payouts of both the long and short legs (puts have maximum payout with spot at zero) is A×B. This can be seen in the figure above on the right (showing a 99-101 101x99 ratio put spread). Looking at such leveraged ratio put spreads enforces skew decaying by time, not by the square root of time. However, for reasonable values of skew this condition only applies for long maturities (c5 years).

PROOF SKEW IS CAPPED AND FLOORED BY SQUARE ROOT OF TIME

Enforcing positive values for put and call spreads is the same as the below two conditions:

- **Change in price of a call when strike increases has to be negative** (intuitively makes sense, as you have to pay more to exercise the higher strike call).
- **Change in price of a put when strike increases has to be positive** (intuitively makes sense, as you receive more value if the put is exercised against you).

These conditions are the same as saying the gradient of x (=Strike/Forward) is bound by:

$$\text{Lower bound} = -\sqrt{2\pi}.e^{\frac{d_1^2}{2}}\left[1-N(d_2)\right] \leq \text{x} \leq \sqrt{2\pi}.e^{\frac{d_1^2}{2}}.N(d_2) = \text{upper bound}$$

It can be shown that these bounds decay by (roughly) the square root of time. This is plotted below.

Figure 139. Upper and Lower Bound for Skew (given 25% volatility)

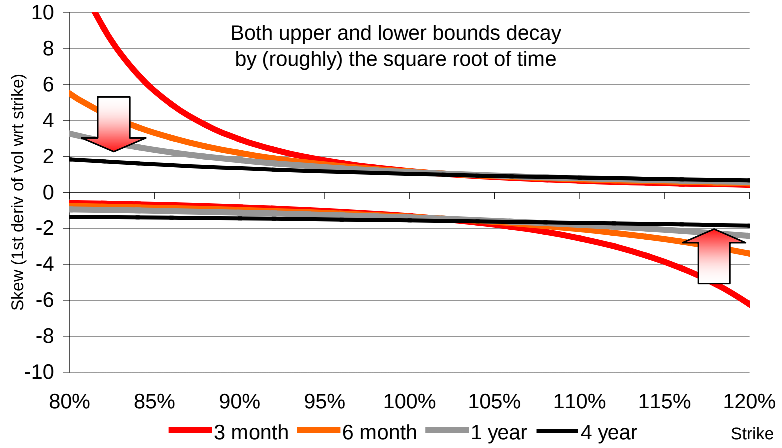

Source: Santander Investment Bolsa.

Proof of theoretical cap for skew works in practice

In the above example, for a volatility of 25% the mathematical lower bound for one-year skew (gradient of volatility with respect to strike) is -1.39. This is the same as saying that the maximum difference between 99% and 100% strike implied is 1.39% (ie, 90%-100% or 95%-105% skew is capped at 13.9%). This theoretical result can be checked by pricing one-year put options with Black-Scholes.

- Price 100% put with 25% implied = 9.95%
- Price 99% put with 26.39% implied = 9.95% (difference of implied of 1.39%)

In practice, skew is likely to be bounded well before mathematical limits

While a 90%-100% one-year skew of 13.9% is very high for skew, we note buying cheap put spreads will appear to be attractive long before the price is negative. Hence, in practice, traders are likely to sell skew long before it hits the mathematical bounds for arbitrage (as a put spread's price tends to zero as skew approaches the mathematical bound). However, as the mathematical bound decays by the square root of time, so too should the 'market bound'.

OTM IMPLIEDS AT THE WINGS HAVE TO BE FLAT IN LOG SPACE

While it is popular to plot implieds vs delta, it can be shown for many models[50] that implied volatility must be linear in log strike (ie, Ln[K/F]) as log strike goes to infinity. Hence a parameterisation of a volatility surface should, in theory, be parameterised in terms of log strike, not delta. In practice, however, as the time value of options for a very high strike is very small, modelling implieds against delta can be used as the bid-offer should eliminate any potential arbitrage.

[50] Eg, stochastic volatility plus jump models.

BLACK-SCHOLES FORMULA

The most popular method of valuing options is the Black-Scholes-Merton model. We show the key formulas involved in this calculation. The assumptions behind the model are also discussed.

BLACK-SCHOLES MAKES A NUMBER OF ASSUMPTIONS

It is often joked that Black-Scholes is the wrong model with the wrong assumptions that gets the right price. The simplicity of the model has ensured that it is still used despite the competition from other, more complicated models. The assumptions are below:

- Constant (known) volatility
- Constant interest rates
- No dividends (a constant dividend yield can, however, be incorporated into the interest rate)
- Zero borrow cost, zero trading cost and zero taxes
- Constant trading
- Stock price return is log normally distributed
- Can trade infinitely divisible amounts of securities
- No arbitrage

BLACK-SCHOLES PRICE OF CALL AND PUT OPTIONS

Call option price = $S \times N(d_1) - Ke^{-rT} N(d_2)$

Put option price = $-S \times N(-d_1) + Ke^{-rT} N(d_2)$

where

$$d_1 = \frac{Ln\left(\frac{S}{K}\right) + (r + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$$

$$d_2 = d_1 - \sigma\sqrt{T} = \frac{Ln\left(\frac{S}{K}\right) + (r - \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$$

S = Spot

K = Strike

r = risk free rate (– dividend yield)

σ = volatility

T = time (years)

GREEKS AND THEIR MEANING

Greeks is the name given to the (usually) Greek letters used to measure risk. We give the Black-Scholes formula for the key Greeks and describe which risk they measure.

VEGA IS NOT A GREEK LETTER

Although Vega is a Greek, it is not a Greek letter. It is instead the brightest star in the constellation Lyra. The main greeks and their definition are in the table below.

Figure 140. Greeks and their definition

Greek	Symbol	Measures	Definition
Delta	δ or Δ	Equity exposure	Change in option price due to spot
Gamma	γ or Γ	Convexity of payout	Change in delta due to spot
Theta	θ or Θ	Time decay	Change in option price due to time passing
Vega	ν	Volatility exposure	Change in option price due to volatility
Rho	ω or Ω	Interest rate exposure	Change in option price due to interest rates
Volga	λ or Λ	Vol of vol exposure	Change in vega due to volatility
Vanna	ψ or Ψ	Skew	Change in vega due to spot OR change in delta due to volatility

Source: Santander Investment Bolsa estimates.

The variables for the below formulae are identical to the earlier definitions in the previous section *Black-Scholes Formula*. In addition:

N'(z) is the normal density function, $\frac{e^{\frac{-x^2}{2}}}{\sqrt{2\pi}}$

N(z) is the cumulative normal distribution, ie, N(0) = 0.5.

DELTA MEASURES EQUITY EXPOSURE

Call delta:

$N(d_1)$

Put delta:

$-N(-d_1)$

= $N(d_1)-1$

The most commonly examined Greek is delta, as it gives the equity sensitivity of the option (change of option price due to change in underlying price). Delta is normally quoted in percent. For calls it lies between 0% (no equity sensitivity) and 100% (trades like a stock). The delta of puts lies between -100% (trades like short stock) and 0%. If a call option has a delta of 50% and the underlying rises €1, the call option increases in value €0.50 (= €1 * 50%). Note the values of the call and put delta in the formula below give the equity sensitivity of a forward of the same maturity as the option expiry. The equity sensitivity to spot is slightly different. Please note that there is a (small) difference between the probability that an option expires ITM and delta.

Figure 141. Delta (for Call) **Delta (for Put)**

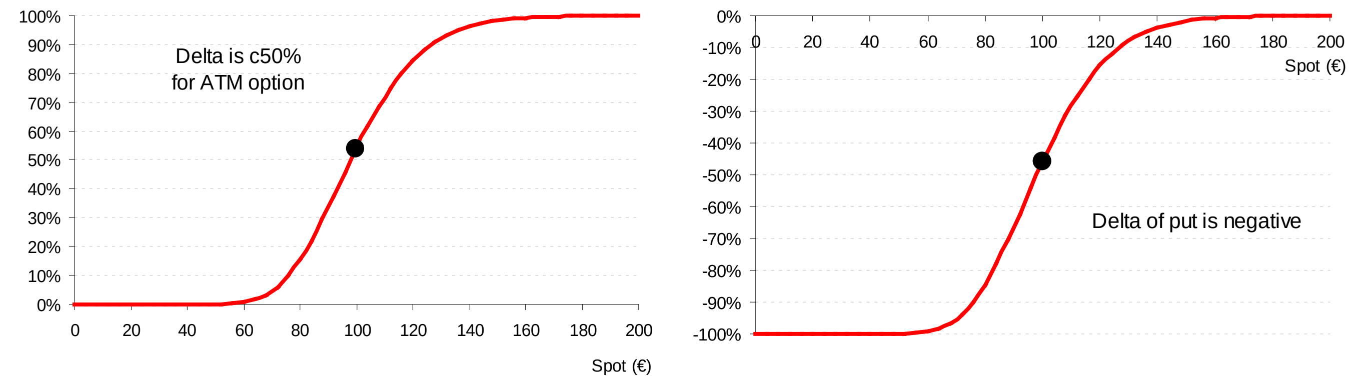

Source: Santander Investment Bolsa estimates.

GAMMA MEASURES CONVEXITY (AMOUNT EARNED DELTA HEDGING)

Gamma:

$$-\frac{N'(d_1)}{S\sqrt{T}\sigma}$$

Gamma measures the change in delta due to the change in underlying price. The higher the gamma, the more convex is the theoretical payout. Gamma should not be considered a measure of value (low or high gamma does not mean the option is expensive or cheap); implied volatility is the measure of an option's value. Options are most convex, and hence have the highest gamma, when they are ATM and also about to expire. This can be seen intuitively as the delta of an option on the day of expiry will change from c0% if spot is just below expiry to c100% if spot is just above expiry (a small change in spot causes a large change in delta; hence, the gamma is very high).

Figure 142. Gamma **Theta**

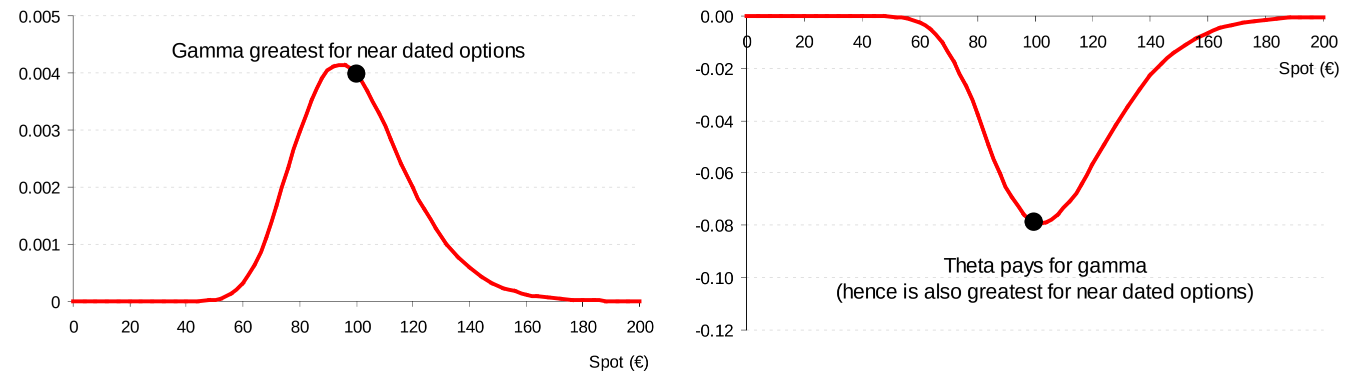

Source: Santander Investment Bolsa estimates.

Call theta:

$$-\frac{S\sigma \times N'(d_1)}{2\sqrt{T}}$$

$$-rKe^{-rT}N(d_2)$$

Put theta:

$$-\frac{S\sigma \times N'(d_1)}{2\sqrt{T}}$$

$$+rKe^{-rT}N(-d_2)$$

THETA MEASURES TIME DECAY (COST OF BEING LONG GAMMA)

Theta is the change in the price of an option for a change in time to maturity; hence, it measures time decay. In order to find the daily impact of the passage of time, its value is normally divided by 252 (trading days in the year). If the second term in the formula below is ignored, the theta for calls and puts are identical and proportional to gamma. Theta can therefore be considered the cost of being long gamma.

VEGA MEASURES VOLATILITY EXPOSURE (AVERAGE OF GAMMAS)

Vega:
$S\sqrt{T} \times N'(d_1)$

Vega gives the sensitivity to volatility of the option price. Vega is normally divided by 100 to give the price change for a 1 volatility point (ie, 1%) move in implied volatility. Vega can be considered to be the average gamma (or non-linearity) over the life of the option. As vega has a √T in the formula power vega (vega divided by square root of time) is often used as a risk measure (to compensate for the fact that near dated implieds move more than far-dated implieds).

Figure 143. Vega

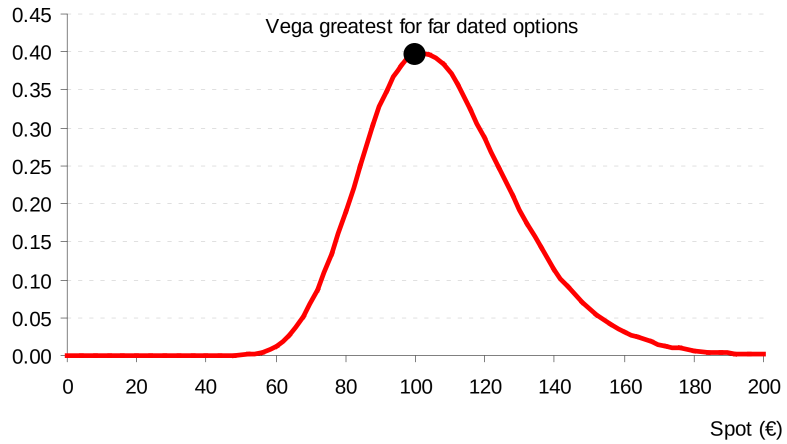

Source: Santander Investment Bolsa.

RHO MEASURES INTEREST RATE RISK (RELATIVELY SMALL)

Call rho:
$KTe^{-rT}N(d_2)$
Put rho:
$-KTe^{-rT}N(-d_2)$

Rho measures the change in the value of the option due to a move in the risk-free rate. The profile of rho vs spot is similar to delta, as the risk-free rate is more important for more equity-sensitive options (as these are the options where there is the most benefit in selling stock and replacing it with an option and putting the difference in value on deposit). Rho is normally divided by 10,000 to give the change in price for a 1bp move.

Figure 144. Rho (for call)

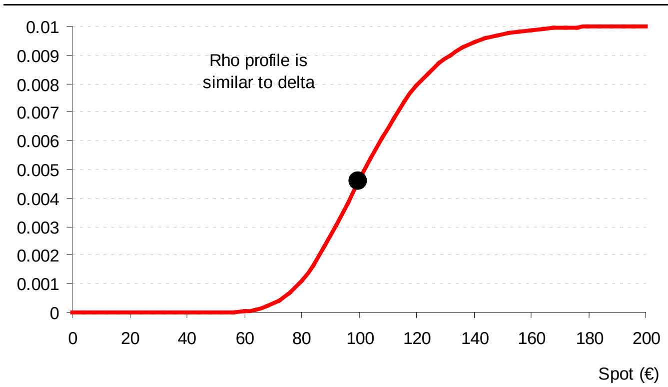

Rho (for put)

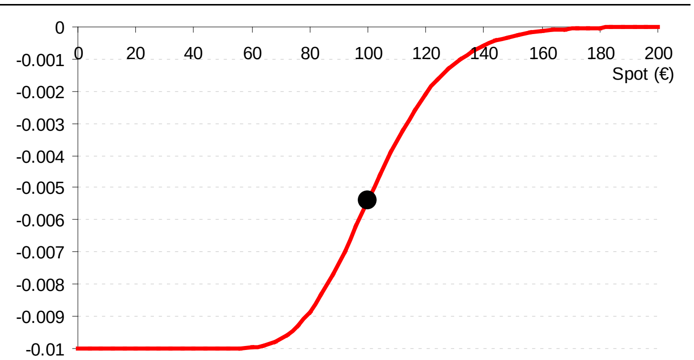

Source: Santander Investment Bolsa estimates.

VOLGA MEASURES VOLATILITY OF VOLATILITY EXPOSURE

Volga:

$$\frac{S\sqrt{T}d_1 d_2 N'(d_1)}{\sigma}$$

Volga is short for VOLatility GAmma, and is the rate of change of vega due to a change in volatility. Volga (or Vomma/vega convexity) is highest for OTM options (approximately 10% delta), as these are the options where the probability of moving from OTM to ITM has the greatest effect on its value. For more detail on Volga, see the section *How to Measure Skew and Smile*.

Figure 145. Volga

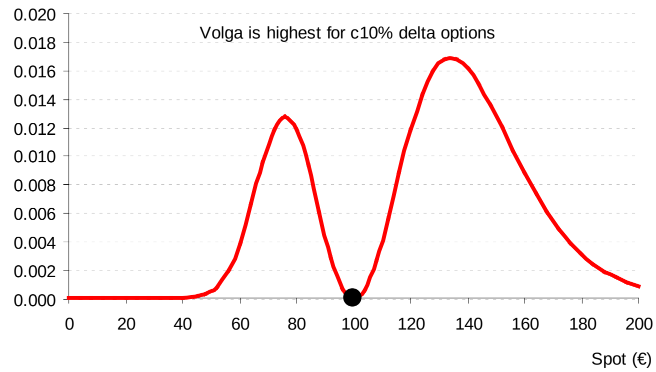

Vanna

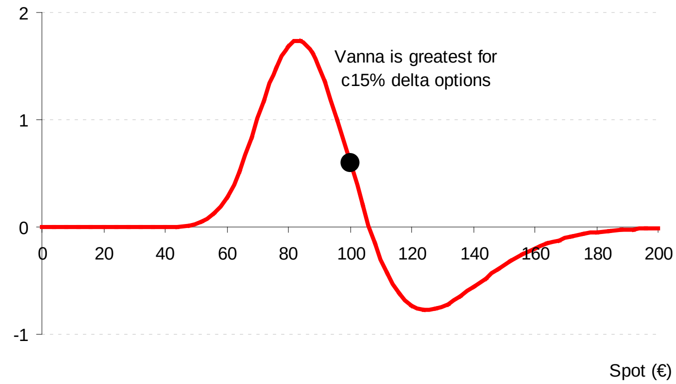

Source: Santander Investment Bolsa estimates.

VANNA MEASURES SKEW EXPOSURE

Vanna:

$$\frac{-d_2 N'(d_1)}{\sigma}$$

Vanna has two definitions as it measures the change in vega given a change in spot and the change in delta due to a change in volatility. The change in vega for a change in spot can be considered to measure the skew position, as this will lead to profits on a long skew trade if there is an increase in volatility as spot declines. The extreme values for vanna occur for c15 delta options, similar to volga's c10 delta peaks. For more detail on vanna, see the section *How to Measure Skew and Smile*.

ADVANCED (PRACTICAL OR SHADOW) GREEKS

How a volatility surface changes over time can impact the profitability of a position. While the most important aspects have already been covered (and are relatively well understood by the market) there are 'second order' Greeks that are less well understood. Two of the most important are the impact of the passage of time on skew (volatility slide theta), and the impact of a movement in spot on OTM options (anchor delta). These Greeks are not mathematical Greeks, but are practical or 'shadow' Greeks.

INCREASE IN SKEW AS TIME PASSES CAUSES 'VOL SLIDE THETA'

As an option approaches expiry, its maturity decreases. As near-dated skew is larger than far-dated skew, the skew of a fixed maturity option will increase as time passes. This can be seen by assuming that skew by maturity (eg, three-month or one-year) is constant (ie, relative time, the maturity equivalent of sticky moneyness or sticky delta). We also assume that three-month skew is larger than the value of one year skew. If we buy a low strike one year option (ie, we are long skew) then, assuming spot and ATM volatility stay constant, when the option becomes a three-month option its implied will have risen (as three-month skew is larger than one-year skew and ATM volatility has not changed). We define 'volatility slide theta' as the change in price of an option due to skew increasing with the passage of time[51].

VOLATILITY SLIDE THETA IS MOST IMPORTANT FOR NEAR EXPIRIES

Volatility slide theta measures the increase of skew as expiry approaches

Given that skew increases as maturity decreases, this change in skew will increase the value of long skew positions (as in the example) and decrease the value of short skew positions. The effect of 'volatility slide theta' is negligible for medium- to far-dated maturities, but increases in importance as options approach expiry. If a volatility surface model does not take into account 'volatility slide theta', then its impact will be seen when a trader re-marks the volatility surface.

VOL SLIDE THETA MEASURES IMPACT OF CONSTANT SMILE RULE

The constant smile rule (CSR) details how forward starting options should be priced. The impact of this rule on valuations is given by the 'volatility slide theta' as they both assume a fixed maturity smile is constant. The impact of this assumption is more important for forward starting options than for vanilla options.

[51] While we concentrate on Black-Scholes implied volatilities, volatility slide theta also affects local volatility surfaces.

WHEN TRADERS CHANGE THEIR 'ANCHOR' THIS INTRODUCES A **SECOND ORDER DELTA ('ANCHOR DELTA')**

Volatility surfaces are normally modelled via a parameterisation. One of the more popular parameterisations is to set the ATMf volatility from a certain level of spot, or 'anchor', and then define the skew (slope). While this builds a reasonable volatility surface for near ATM options, the wings will normally need to be slightly adjusted. Normally a fixed skew for both downside puts and upside calls will cause upside calls to be too cheap (as volatility will be floored) and downside puts to be too expensive (as volatility should be capped at some level, even for very low strikes). As the 'anchor' is raised, the implied volatility of OTM options declines (assuming the wing parameters for the volatility surface stay the same). We call this effect 'anchor delta'.

Implied volatility has to be floored, and capped, for values to be realistic

When traders change their 'anchor' this introduces a second order 'anchor delta'

There are many different ways a volatility surface parameterisation can let traders correct the wings, but the effect is usually similar. We shall simply assume that the very OTM call implied volatility is lifted by a call accelerator, and very OTM put implied volatility is lowered by a put decelerator. This is necessary to prevent call implieds going too low (ie, below minimum realised volatility), or put implieds going too high (ie, above maximum realised volatility). The effect of these wing parameters is shown in Figure 146 below.

Figure 146. Skew with Put Decelerators and Call Accelerators

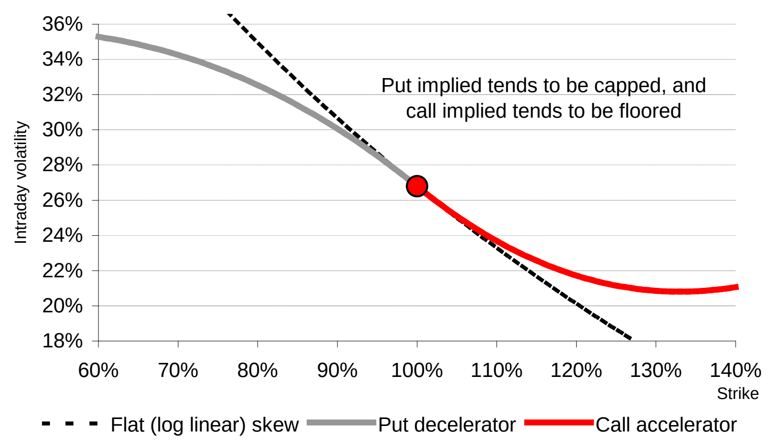

Source: Santander Investment Bolsa.

Traders tend to refresh a surface by only changing the key parameters

For liquid underlyings such as indices, a volatility surface is likely to be updated several times a day (especially if markets are moving significantly). Usually only the key parameters will be changed, and the less key parameters such as the wing parameters are changed less frequently. We shall assume that there will be many occasions where there is a movement in spot along the skew (ie, static strike for near ATM strikes). In these cases a trader is likely to change the anchor (and volatility at the anchor, which has moved along the skew), but leave the remaining skew and wing parameters (which are defined relative to the anchor) unchanged. In order to have the same implied volatility for OTM options after changing the anchor, the call accelerator should be increased and the put decelerator decreased. In practice this does not always happen, as wing parameters are typically changed less frequently. The effect of an increase in anchor along the (static strike) skew while leaving the wing parameters unchanged is shown below.

Figure 147. Moving Anchor 10% Higher Along the Skew

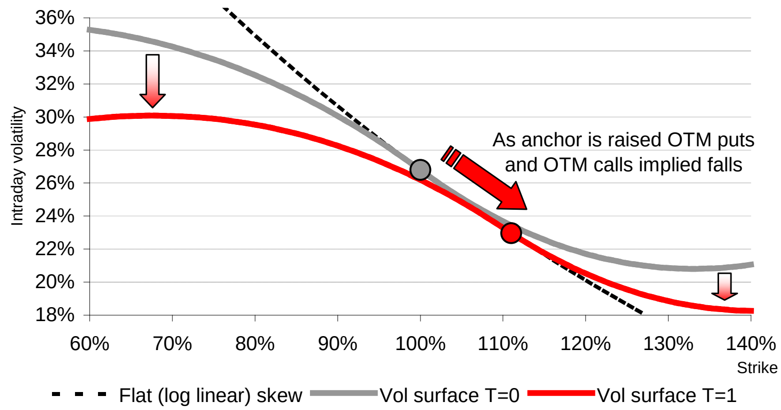

Source: Santander Investment Bolsa.

OTM options have a second order 'anchor delta'

To simplify the example we shall assume the call wing parameter increases the implied volatility for strikes 110% and more, and the put wing parameter decreases the implied volatility for strikes 90% or lower. If spot rises 10%, the 120% call implied volatility will suffer when the anchor is re-marked 10% higher, because the call implied volatility is initially lifted by the call wing parameter (which no longer has an effect). OTM calls therefore have a negative 'anchor delta' as they lose value as anchor rises. Similarly, as anchor rises the effect of the put wing will increase, lowering the implied volatility of puts of strike 90% or less as anchor rises. So, under this scenario all options that are OTM have a negative 'anchor delta' that needs to be hedged.

SHORTING STOCK BY BORROWING SHARES

The hedging of equity derivatives assumes you can short shares by borrowing them. We show the processes involved in this operation. The disadvantages, and benefits, to an investor who lends out shares are also explained.

THERE IS NO COUNTERPARTY RISK WHEN YOU LEND SHARES

Voting rights lost when shares are lent

To short shares initially, the shares must first be borrowed. In order to remove counterparty risk, when an investor lends out shares he/she receives collateral (cash, stock, bonds, etc) for the same value. Both sides retain the beneficial ownership of both the lent security and the collateral, so any dividends, coupons, rights issues are passed between the two parties. If cash is used as collateral, the interest on the cash is returned. Should a decision have to be made, ie, to receive a dividend in cash or stock, the decision is made by the original owner of the security. The only exception is that the voting rights are lost, which is why lent securities are often called back before votes. To ensure there is no counterparty risk during the time the security is lent out, the collateral and lent security is marked to market and the difference settled for cash (while a wide range of securities can be used as initial collateral, only cash can be used for the change in value of the lent security).

Figure 148. Borrowing Shares

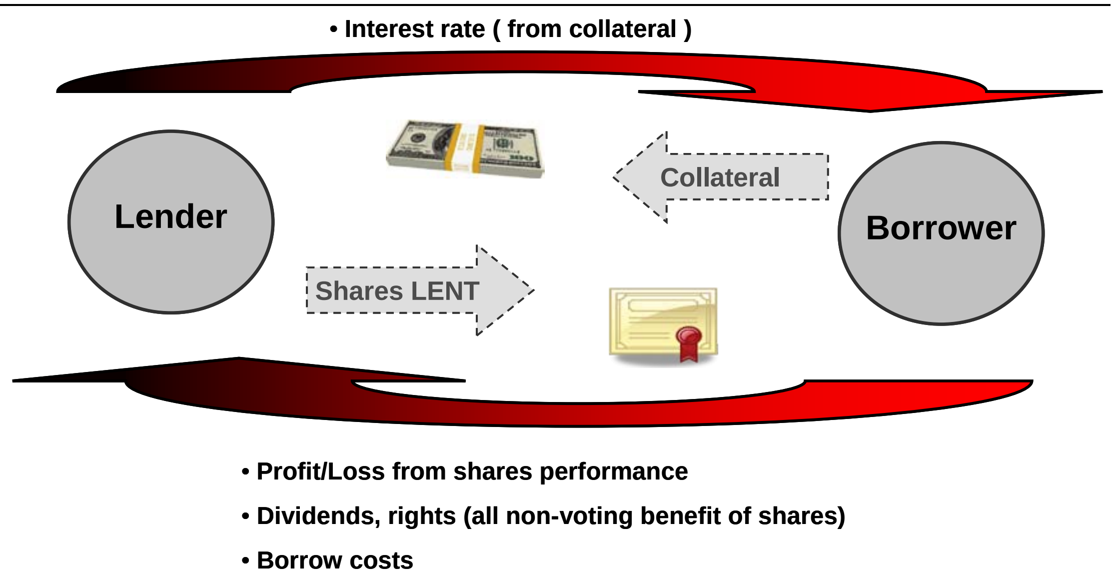

LENDER:

Retains

• Full exposure to movement share price

•All non-voting benefit of shares (dividends, rights)

Gains

• Borrow costs

Loses

• Voting rights of shares

BORROWER:

Retains

• Benefits of collateral

Gains

• Voting rights

• Right to sell shares (go short)

Loses

• Borrow cost

Source: Santander Investment Bolsa.

SELLING THE STOCK YOU HAVE BORROWED GIVES A SHORT POSITION

Once an investor has borrowed shares, these shares can be sold in the market. The proceeds from this sale can be used as the collateral given to the lender. Selling borrowed shares gives a short position, as profits are earned if the stock falls (as it can be bought back at a lower price than it was sold for, and then returned to the original owner).

Figure 149. Borrowing Shares – can you resize so double line is closer

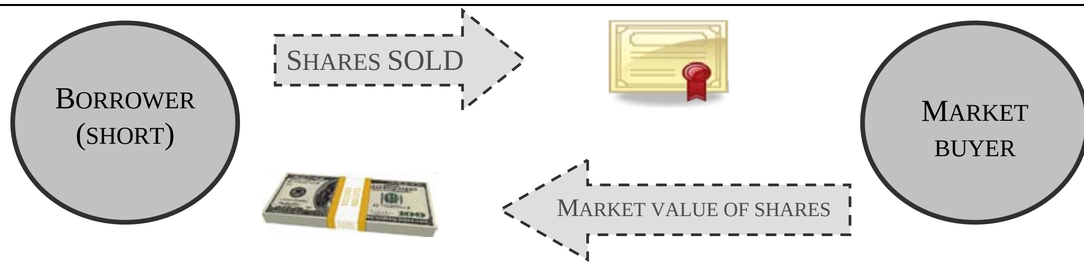

BORROWER:

Retains

- Benefits of collateral
- Responsibility to return shares to lender

Gains

- Profits from share price declines
- Cash value of shares (can be used as collateral)

Loses

- Borrow cost

Source: Santander Investment Bolsa.

EARN SHORT REBATE WHEN YOU SHORT

Short seller receives interest rate on collateral (less dividends and borrow) as short rebate

The investor who has shorted the shares receives interest on the collateral, but has to pass dividends and borrow cost to the original owner. The net of these cash flows is called the short rebate, as it is the profit (or loss for high dividend paying stocks) that occurs if there is no change in the price of the shorted security. Shorting shares is therefore still profitable if shares rise by less than the short rebate.

Short rebate = interest rate (normally central bank risk free rate) – dividends – borrow cost

Figure 150. Initial and Final Position of Lender, Borrower and Market Following Shorting of Shares

	Before lending: Lender	Before lending: Borrow	Before lending: Market		After lending: Lender	After lending: Borrow	After lending: Market	
Equity								
Holds shares	+			⇨			+	Share passes from lender to market (as receives collateral there is no credit risk)
P&L from share movement	+			⇨	+	-	+	Lender still has equity exposure, borrower is short to market
Economic benefit of shares (dividends, rights)	+			⇨	+	-	+	Borrower is short dividends and rights to lender
Voting rights	+			⇨			+	Lender loses voting rights (only difference to owning equity)
Borrow cost				⇨	+	-		Lender earns borrrow cost (compensates loss of voting rights)
Cash (or collateral)								
Cash (or collateral) value of shares			+	⇨	+			Borrower passes value of share sale to lender as collateral (compensates for value of shares borrowed)
Interest on cash value of shares (or benefits of collateral, e.g. coupons)			+	⇨		+		Borrower gets interest on collateral (less dividends and borrow) as short rebate

Source: Santander Investment Bolsa.

SORTINO RATIO

If an underlying is distributed normally, standard deviation is the perfect measure of risk. For returns with a skewed distribution, such as with option trading or call overwriting, there is no one perfect measure of risk; hence, several measures of risk should be used. The Sortino is one of the most popular measures of skewed risk, as it only takes into account downside risk.

SORTINO RATIO IS MODIFICATION OF SHARPE RATIO

Sortino ratio is Sharpe ratio where excess return is divided by downside risk (not total risk)

The Sharpe ratio measures the excess return, or amount of return (R) that is greater than the target rate of return (T). Often zero or risk-free rate is used as the target return. To take volatility of returns into account, this excess return is divided by the standard deviation. However, this takes into account both upside and downside risks. As investors are typically more focused on downside risks, the Sortino ratio modifies the Sharpe ratio calculation to only divide by the downside risk (DR). The downside risk is the square root of the target semivariance, which can be thought of as the amount of standard deviation due to returns less than the target return. The Sortino ratio therefore only penalises large downside moves, and is often thought of as a better measure of risk for skewed returns.

$$Sortino = \frac{R - T}{DR}$$

where

$$DR = \left[\int_{\infty}^{T} (T - x) f(x) dx \right]^{0.5}$$

R = Return

T = Target return

CAPITAL STRUCTURE ARBITRAGE

When Credit Default Swaps were created in the late 1990's, they traded independently of the equity derivative market. The high levels of volatility and credit spreads during the bursting of the TMT bubble demonstrated the link between credit spreads, equity, and implied volatility. We examine four models that demonstrate this link (Merton model, jump diffusion, put vs CDS, and implied no-default volatility).

NORMALLY TRADE CREDIT VS EQUITY, NOT VOLATILITY

Capital structure arbitrage models can link the price of equity, credit and implied volatility. However, the relatively wide bid-offer spreads of equity derivatives mean trades are normally carried out between credit and equity (or between different subordinations of credit and preferred shares vs ordinary shares). The typical trade is for an investor to go long the security that is highest in the capital structure, for example, a corporate bond (or potentially a convertible bond), and short a security that is lower in the capital structure, for example, equity. Reverse trades are possible, for example, owning a subordinated higher yielding bond and shorting a senior lower yielding bond (and earning the positive carry as long as bankruptcy does not occur). Only for very wide credit spreads and high implied volatility is there a sufficiently attractive opportunity to carry about an arbitrage between credit and implied volatility. We shall concentrate on trading credit vs equity, as this is the most common type of trade.

Credit spread is only partly due to default risk

The OAS (Option Adjusted Spread) of a bond over the risk-free rate can be divided into three categories. There is the expected loss from default; however, there is also a portion due to general market risk premium and additionally a liquidity cost. Tax effects can also have an effect on the corporate bond market. Unless a capital structure arbitrage model takes into account the fact that not all of a bond's credit spread is due to the risk of default, the model is likely to fail. The fact that credit spreads are higher than they should be if bankruptcy risk was the sole risk of a bond was often a reason why long credit short equity trades have historically been more popular than the reverse (in addition to the preference to being long the security that is highest in the capital structure in order to reduce losses in bankruptcy).

CDS usually better than bonds for credit leg, as they are unfunded and easier to short

Using CDS is likely to be superior to corporate bonds for funding and technical reasons

Using CDS rather than corporate bonds can reduce many of the discrepancies in spread that a corporate bond suffers and narrow the difference between the estimated credit spread and the actual credit spread. We note that CDS are an unfunded trade (ie, leveraged), whereas corporate bonds are a funded trade (have to fund the purchase of the bond) that has many advantages when there is a funding squeeze (as occurred during the credit crunch). CDS also allow a short position to be easily taken, as borrow for corporate bonds is not always available, is usually expensive and can be recalled at any time. While borrow for bonds was c50bp before the credit crunch it soared to c5% following the crisis.

Credit derivatives do not have established rules for equity events

While credit derivatives have significant language against credit events, they have no language for equity events, such as special dividends or rights issues. Even for events such as takeovers and mergers, where there might be relevant documentation, credit derivatives are likely to behave differently than equity (and equity derivatives).

CREDIT MARKET CAN LEAD EQUITY MARKET AND VICE VERSA

While the bond market lags the CDS market, on average there is no difference between CDS and equity

We note that there are occasions when corporate bond prices lag a movement in equity prices, simply as traders have not always updated levels (but this price would be updated should an investor request a firm price). CDS prices suffer less from this effect, and we note for many large companies the corporate bond market is driven by the CDS market and not vice versa (the tail wags the dog). Although intuitively the equity market should be more likely to lead the CDS market than the reverse (due to high frequency traders and the greater number of market participants), when the CDS market is compared to the equity market on average neither consistently leads the other. Even if the CDS and equity on average react equally as quickly to new news, there are still occasions when credit leads equity and vice versa. Capital structure arbitrage could therefore be used on those occasions when one market has a delayed reaction to new news compared to the other.

GREATEST OPPORTUNITY ON BBB OR BB RATED COMPANIES

In order for capital structure arbitrage to work, there needs to be a strong correlation between credit and equity. This is normally found in companies that are rated BBB or BB. The credit spread for companies with ratings of A or above is normally more correlated to the general credit supply and interest rates than the equity price. For very speculative companies (rated B or below), the performance of their debt and equity is normally very name-specific, and often determined by the probability of takeover or default.

Capital structure arbitrage works best when companies don't default

Capital structure arbitrage is a bet on the convergence of equity and credit markets. It has the best result when a company in financial distress recovers, and the different securities it has issued converge. Should the company enter bankruptcy, the returns are less impressive. The risk to the trade is that the company becomes more distressed, and as the likelihood of bankruptcy increases the equity and credit markets cease to function properly. This could result in a further divergence or perhaps closure of one of the markets, potentially forcing a liquidation of the convergence strategy.

FUNDAMENTAL FACTORS CAN DWARF STATISTICAL RELATIONSHIPS

Capital structure arbitrage assumes equity and credit markets move in parallel. However, there are many events that are bullish for one class of investors and bearish for another. This normally happens when the leverage of a company changes suddenly. Takeovers and rights issues are the two main events that can quickly change leverage. Special dividends, share buybacks and a general reduction of leverage normally have a smaller, more gradual effect.

Rights issue. A rights issue will always reduce leverage, and is effectively a transfer of value from equity holders to debt holders (as the company is less risky, and earnings are now divided amongst a larger number of shares).

Takeover bid (which increases leverage). When a company is taken over, unless the acquisition is solely for equity, a portion of the acquisition will have to be financed with cash or debt (particularly during a leveraged buyout). In this case, the leverage of the acquiring company will increase, causing an increase in credit spreads and a reduction in the value of debt. Conversely, the equity price of the acquiring company is more stable. For the acquired company, the equity price should jump close to the level of the bid and, depending on the structure of the offer, the debt could fall (we note that if the acquired company is already in distress the value of debt can rise; for example, when Household was acquired by HSBC).

GM EQUITY SOARED A DAY BEFORE CREDIT SANK, CAUSING LOSSES

GM was downgraded by S&P a day after Kirk Kerkorian bought a 5% stake

On May 4, 2005, Kirk Kerkorian announced the intention to increase his (previously unknown) stake in GM, causing the troubled company's share price to soar 18% intraday (7.3% close to close). The following day, S&P downgraded GM and Ford to 'junk', causing a collapse in the credit market and a 122bp CDS rise in two days. As many capital structure arbitrage investors had a long credit short equity position, both legs were loss making and large losses were suffered.

Figure 151. Equity vs Credit Spread (5-Year CDS) for Ford

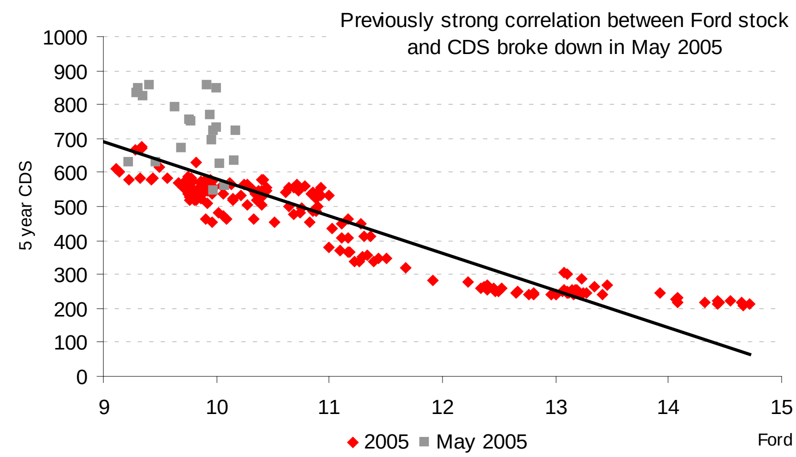

Source: Santander Investment Bolsa.

CORRELATION BETWEEN CREDIT AND EQUITY LOW AT STOCK LEVEL

For many companies the correlation between equity and credit is not particularly strong, with a typical correlation between 5% and 15%. Hence it is necessary for a capital structure arbitrage investor to have many different trades on simultaneously. The correlation of a portfolio of bonds and equities is far higher (c90%).

Figure 152. SX5E vs 5-Year CDS (European HiVol)

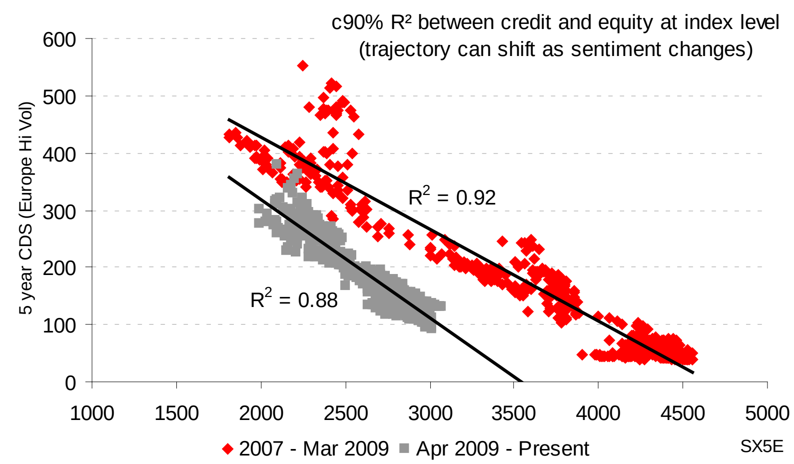

Source: Santander Investment Bolsa.

MODELLING THE LINK BETWEEN CREDIT SPREADS AND IMPLIED VOL

While there are many models that show the link between the equity, equity volatility and debt of a company, we shall restrict ourselves to four of the most popular.

- **Merton model.** The Merton model uses the same model as Black-Scholes, but applies it to a firm. If a firm is assumed to have only one maturity of debt, then the equity of the company can be considered to be a European call option on the value of the enterprise (value of enterprise = value of debt + value of equity) whose strike is the face value of debt. This model shows how the volatility of equity rises as leverage rises. The Merton model also shows that an increase in volatility of the enterprise increases the value of equity (as equity is effectively long a call on the value of the enterprise), and decreases the value of the debt (as debt is effectively short a put on the enterprise, as they suffer the downside should the firm enter bankruptcy but the upside is capped).

- **Jump diffusion.** A jump diffusion model assumes there are two parts to the volatility of a stock. There is the diffusive (no-default) volatility, which is the volatility of the equity without any bankruptcy risk, and a separate volatility due to the risk of a jump to bankruptcy. The total volatility is the sum of these two parts. While the diffusive volatility is constant, the effect on volatility due to the jump to bankruptcy is greater for options of low strike than high strike causing 'credit induced skew'. This means that as the credit spread of a company rises, this increases the likelihood of a jump to bankruptcy and increases the skew. A jump diffusion model therefore shows a link between credit spread and implied volatility.

- **Put vs CDS.** As the share price of a company in default tends to trade close to zero, a put can be assumed to pay out its strike in the even of default. This payout can be compared to the values of a company's CDS, or its debt market (as the probability of a default can be estimated from both). As a put can also have a positive value even if a company does not default, the value of a CDS gives a floor to the value of puts. As 1xN put spreads can be constructed to never have a negative payout, various caps to the value of puts can be calculated by ensuring they have a cost. The combination of the CDS price floor, and put price cap, gives a channel for implieds to trade without any arbitrage between CDS and put options.

- **No-default implied volatility.** Using the above put vs CDS methodology, the value of a put price due to the payout in default can be estimated. If this value is taken away from the put price, the remaining price can be used to calculate a no-default implied volatility (or implied diffusive volatility). The skew and term structure of implied no-default implied volatilities are flatter than Black-Scholes implied volatility, which allows an easier comparison and potential for identifying opportunities.

(1) MERTON MODEL

Value of enterprise is equal to sum of equity and debt

The Merton model assumes that a company has an enterprise value (V) whose debt (D) consists of only one zero coupon bond whose value at maturity is K. These assumptions are made in order to avoid the possibility of a default before maturity (which would be possible if there was more than one maturity of debt, or a coupon had to be paid. The company has one class of equity (E) that does not pay a dividend. The value of equity (E) and debt (D) at maturity is given below.

Enterprise value = V = E + D

Equity = Max(V – K, 0) = call on V with strike K

Debt = Min(V, K) = K – Max(K – V, 0) = Face value of debt K – put on V with strike K

Figure 153. Graph of Value of Enterprise, Equity and Debt

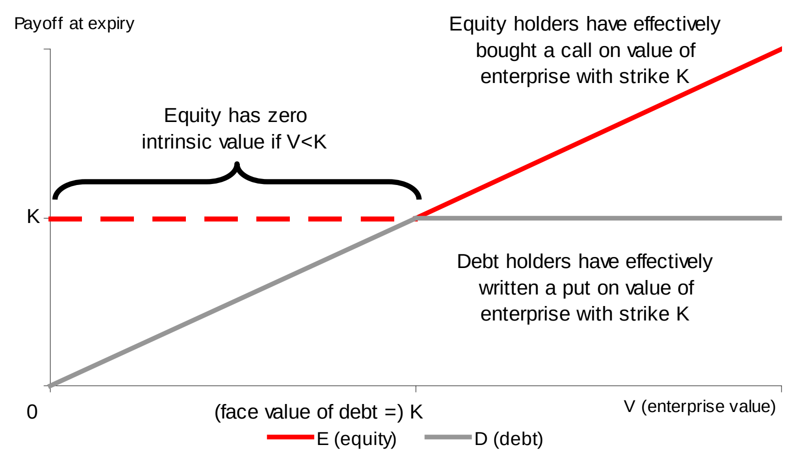

Source: Santander Investment Bolsa.

Enterprise value of a firm at maturity has to be at least K or it will enter bankruptcy

Before the maturity of the debt, the enterprise has obligations to both the equity and debt holders. At the maturity of the debt, if the value of the enterprise is equal to or above K, the enterprise will pay off the debt K and the remaining value of the firm is solely owned by the equity holders. If the value of the enterprise is below K then the firm enters bankruptcy. In the event of bankruptcy, the equity holders get nothing and the debt holders get the whole value of the enterprise V (which is less than K).

Equity is long a call on the value of a firm

Equity holders are long a call, put holders are short a put

If the value of the enterprise V is below the face value of debt K at maturity the equity holders receive nothing. However, if V is greater than K, the equity holders receive V - K. The equity holders therefore receive a payout equal to a call option on V of strike K.

Debt is short a put on value of firm

The maximum payout for owners of debt is the face value of debt. This maximum payout is reduced by the amount the value of the enterprise is below the face value of debt at maturity. Debt is therefore equal to the face value of debt less the value of a put on V of strike K.

DEBT HAS A DELTA THAT CAN BE USED TO ARBITRAGE VS EQUITY

As the value of the short put has a delta, debt has a delta. It is therefore possible to go long debt and short equity (at the calculated delta using the Merton model) as part of a capital structure arbitrage trade.

If enterprise value is unchanged, then if value of equity rises, value of credit falls

Equity holders are long vol, debt holders are short vol

As enterprise value is equal to the sum of equity and debt, if enterprise value is kept constant then for equity to rise the value of debt must fall. An example would be if a company attempts to move into higher-risk activity, lifting its volatility. As equity holders are long a call on the value of the company they benefit from the additional time value. However, as debt holders are short a put they suffer should a firm move into higher-risk activities.

Merton model assumes too high a recovery rate

Using the vanilla Merton model gives unrealistic results with credit spreads that are too tight. This is because the recovery rate (of V/K) is too high. However, using more advanced models (eg, stochastic barrier to take into account the default point is unknown), the model can be calibrated to market data.

MERTON MODEL EXPLAINS EQUITY SKEW

The volatility of an enterprise should be based on the markets in which it operates, interest rates and other macro risks. It should, however, be independent of how it is funded. The proportion of debt to equity therefore should not change the volatility of the enterprise V; however, it does change the volatility of the equity E. It can be shown that the volatility of equity is approximately equal to the volatility of the enterprise multiplied by the leverage (V/E). Should the value of equity fall, the leverage will rise, lifting the implied volatility. This explains skew: the fact that options of lower strike have an implied volatility greater than options of high strike.

$\sigma_E \approx \sigma_V \times V / E \; (= \sigma_V \times \text{leverage})$

Figure 154. Value of Enterprise and Equity with Low Debt

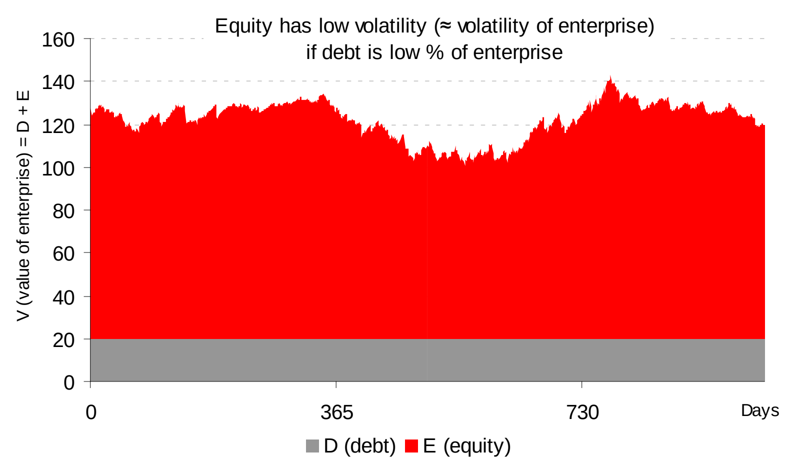

Source: Santander Investment Bolsa.

Firms with a small amount of debt have equity volatility roughly equal to firm volatility

If a firm has a very small (or zero) amount of debt, then the value of equity and the enterprise are very similar. In this case, the volatility of the equity and enterprise should be very similar (see Figure 154 above).

Firms with high value of debt to equity have very high equity volatility

For enterprises with very high levels of debt, a relatively small percentage change in the value of the enterprise V represents a relatively large percentage change in the value of equity. In these cases equity volatility will be substantially higher than the enterprise volatility (see Figure 155 below).

Figure 155. Value of Enterprise and Equity with High Debt

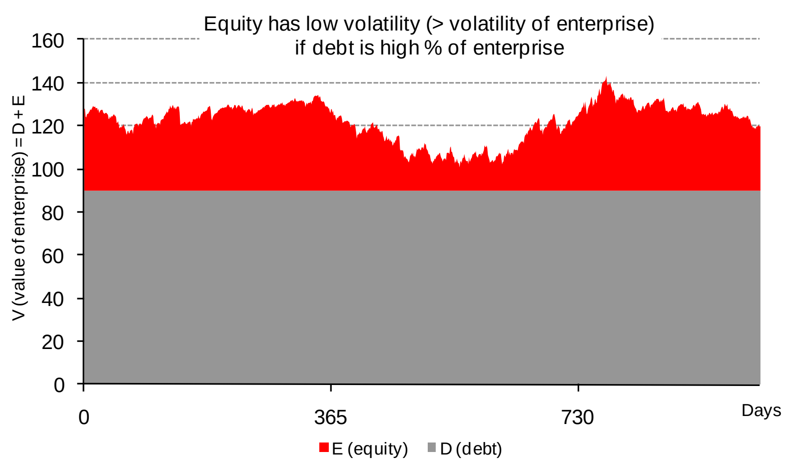

Source: Santander Investment Bolsa.

Proof equity volatility is proportional to leverage

The mathematical relationship between the volatility of the enterprise and volatility of equity is given below. The $N(d_1)$ term adjusts for the delta of the equity.

$$\sigma_E = N(d_1) \times \sigma_V \times V / E$$

Equity volatility is proportional to enterprise volatility multiplied by leverage

If we assume the enterprise is not distressed and the equity is ITM, then $N(d_1)$ or delta of the equity should be very close to 1 (it is usually c90%). Therefore, the equation can be simplified so the volatility of equity is proportional to leverage (V / E).

$$\sigma_E \approx \sigma_V \times \text{leverage}$$

(2) JUMP DIFFUSION

A jump diffusion model separates the movement of equities into two components. There is the diffusive volatility, which is due to random log-normally distributed returns occurring continuously over time. In addition, there are discrete jumps the likelihood of which is given by a credit spread. The total of the two processes is the total volatility of the underlying. It is this total volatility that should be compared to historic volatility or Black-Scholes volatility.

Default risk explained by credit spread

Zero credit spread implies a company can never default

For simplicity, we shall assume that in a jump diffusion model the jumps are to a zero stock price as the firm enters bankruptcy, but results are similar for other assumptions. The credit spread determines the risk of entering bankruptcy. If a zero credit spread is used, the company will never default. The probability of default increases as the credit spread increases (approximately linearly).

Figure 156. Credit-Induced Skew (with 100bp credit spread)

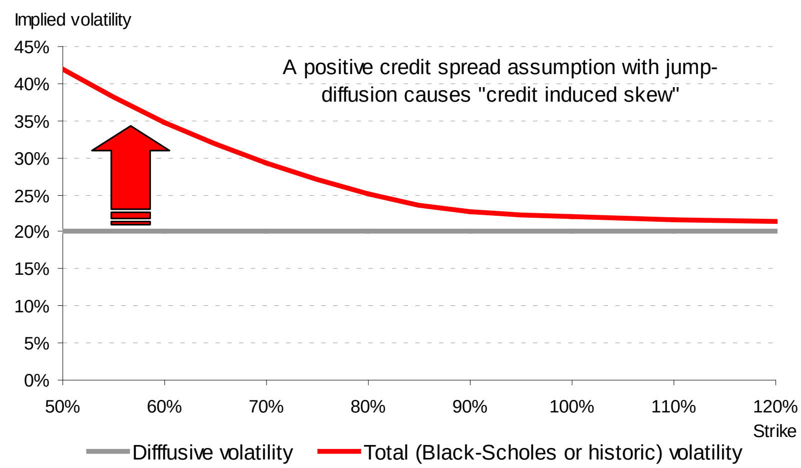

Source: Santander Investment Bolsa.

JUMP DIFFUSION CAUSES CREDIT-INDUCED SKEW

To show how credit spread (or bankruptcy) causes credit-induced skew, we shall price options of different strike with jump diffusion, keeping the diffusive volatility and credit spread constant. Using the price of the option, we shall then calculate the Black-Scholes implied volatility. The Black-Scholes implied volatility is higher for lower strikes than higher strikes, causing skew.

Credit-induced skew is caused by 'option on bankruptcy'

The time value of an option will be divided between the time value due to diffusive volatility and the time value due to the jump to zero in bankruptcy. High strike options will be relatively unaffected by the jump to bankruptcy, and the Black-Scholes implied volatility will roughly be equal to the diffusive volatility. However, the value of a jump to a zero stock price will be relatively large for low strike put options (which, due to put call parity, is the implied for all options). The difference between the Black-Scholes implied and diffusive volatility could be considered to be the value due to the 'option on bankruptcy'.

(3) PUT VS CDS

The probability distribution of a stock price can be decomposed into the probability of a jump close to zero due to credit events or bankruptcy, and the log-normal probability distribution that occurs when a company is not in default. While the value of a put option will be based on the whole probability distribution, the value of a CDS will be driven solely by the probability distribution due to default. The (bi-modal) probability distribution of a stock price due to default, and when not in default, is shown below.

Figure 157. Stock Price Probability Distribution

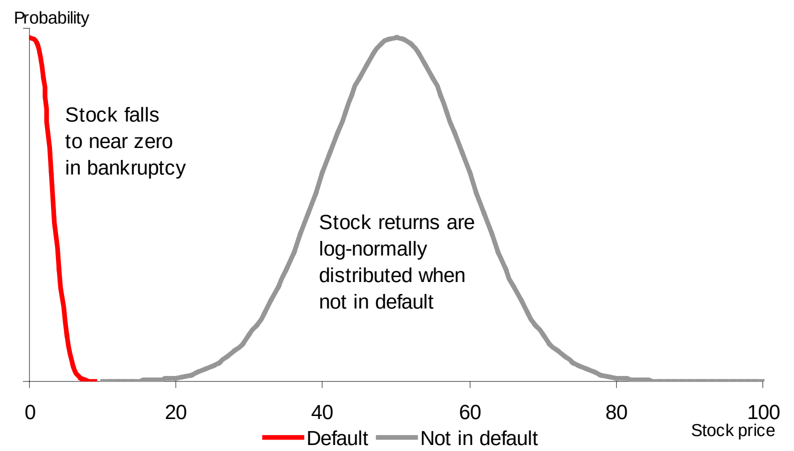

Source: Santander Investment Bolsa.

Puts can be used instead of CDS (as puts pay out strike price in event of bankruptcy)

Stock price jumps to near zero in event of default

When a stock defaults, the share price tends to fall to near zero. The recovery rate of equities can only be above zero if debt recovers 100% of face value, and most investors price in a c40% recovery rate for debt. A put can therefore be assumed to pay out the maximum level (ie, the strike) in the event of default. Puts can therefore be used as a substitute for a CDS. The number of puts needed is shown below.

Value of puts in default = Strike × Number of Puts

Value of CDS in default = (100% – Recovery Rate) × Notional

In order to substitute value of puts in default has to equal value of CDS in default.

⇨ Strike × Number of Puts = (100% – Recovery Rate) × Notional

⇨ Number of Puts = (100% – Recovery Rate) × Notional / Strike

CDS PRICES PROVIDE FLOOR FOR PUTS

Long put vs short CDS was popular in 2000-03 bear market

As a put can have a positive value even if a stock is not in default, a CDS must be cheaper than the equivalent number of puts (equivalent number of puts chosen to have same payout in event of default, ie, using the formula above). If a put is cheaper than a CDS, an investor can initiate a long put-short CDS position and profit from the difference. This was a popular capital structure arbitrage trade in the 2000-03 bear market, as not all volatility traders were as focused on the CDS market as they are now, and arbitrage was possible.

CDS in default must have greater return than put in default (without arbitrage)

As a CDS has a lower price for an identical payout in default, a CDS must have a higher return in default than a put. Given this relationship, it is possible to find the floor for the value of a put. This assumes the price of a CDS is ‘up front’ ie, full cost paid at inception of the contract rather than quarterly.

Puts return in default = Strike / Put Price

CDS return in default = (100% – Recovery Rate) / CDS Price

As CDS return in default must be greater than or equal to put return in default.

- ⇨ (100% – Recovery Rate) / CDS Price ≥ Strike / Put Price
- ⇨ Put Price ≥ Strike × CDS Price / (100% – Recovery Rate)

PUT VS CDS IS A POPULAR CAPITAL STRUCTURE ARBITRAGE TRADE

As the prices of the put and CDS are known, the implied recovery rate can be backed out using the below formula. If an investor’s estimate of recovery value differs significantly from this level, a put vs CDS trade can be initiated. For a low (or zero) recovery rate, the CDS price is too high and a short CDS long put position should be initiated. Conversely, if the recovery rate is too high, a CDS price is too cheap and the reverse (long CDS, short put) trade should be initiated.

Put Price = Strike × CDS Price / (100% – Implied Recovery Rate)

RATIO PUT SPREADS CAP VALUE OF PUTS

CDSs provide a floor to the price of a put. It is also possible to cap the price of a put by considering ratio put spreads. For example, if we have the price for the ATM put, this means we know that the value of a 50% strike put cannot be greater than half the ATM put price. If not, we could purchase an ATM-50% 1×2 put spread (whose payout is always positive) and earn a premium for free. This argument can be used for all strikes K and all 1xN put spreads, and is shown below:

$$N \times \text{put of strike } \frac{K}{N} \leq \text{put of strike } K$$

ARBITRAGE MOST LIKELY WITH LOW STRIKE AND LONG MATURITY

Implied volatility is floored by CDSs, and capped by put ratios

The combination of CDS prices providing a floor, and put prices of higher strikes providing a cap, gives a corridor for the values of puts. The width of this corridor is narrowest for low strike long maturity options, as these options have the greatest percentage of their value associated with default risk. As for all capital structure arbitrage strategies, companies with high credit spreads are more likely to have attractive opportunities and arbitrage is potentially possible for near-dated options.

(4) NO-DEFAULT IMPLIED VOLATILITY

The volatility of a stock price can be decomposed into the volatility due to credit events or bankruptcy and the volatility that occurs when a company is not in default. This is similar to the volatility due to jumps and the diffusive volatility of a jump diffusion model. As the value of a put option due to the probability of default can be calculated from the CDS or credit market, if this value was taken away from put prices this would be the 'no-default put price' (ie, the value the put would have if a company had no credit risk). The implied volatility calculated using this 'no-default put price' would be the 'no-default implied volatility'. No-default implied volatilities are less than the vanilla implied volatility, as vanilla implied volatilities include credit risk).

No-default implied volatilities have lower skew and term structure

While we derive the no-default implied volatility from put options, due to put call parity the implied volatility of calls and puts is identical for European options. As the value of a put associated with a jump to default is highest for low-strike and/or long-dated options, no-default implied volatilities should have a lower skew and term structure than vanilla Black-Scholes implied volatilities. A no-default implied volatility surface should therefore be flatter than the standard implied volatility surface and, hence, could be used to identify potential trading opportunities.

The products and strategies addressed in this report are complex, typically involve a high degree of risk and are intended for sale only to sophisticated investors who are capable of understanding and assuming the risks involved. The market value of any structured security may be affected by changes in economic, financial and political factors (including, but not limited to, spot and forward interest and exchange rates), time to maturity, market conditions and volatility, and the credit quality of any issuer or reference issuer. Any investor interested in purchasing a structured product should conduct their own investigation and analysis of the product and consult with their own professional advisers as to the risks involved in making such a purchase.

The opinions and recommendations included in this report are not necessarily those of the Equity Research Department of Santander Investment Bolsa or of its affiliates. A “Trading Places” rating on a specific company equating to that associated with a conventional “Buy, Hold or Underweight” recommendation should not be construed as a fundamental or official rating of a Santander Investment Bolsa analyst. Furthermore, the opinions and strategies contained in this report are completely independent of those that the Equity Research and Sales/Trading Departments of Santander Investment Bolsa may have from time to time.

Some investments discussed in this report may have a high level of volatility. High volatility investments may experience sudden and large falls in their value causing losses when that investment is realised. Those losses may equal your original investment. Indeed, in the case of some investments the potential losses may exceed the amount of initial investment and, in such circumstances, you may be required to pay more money to support those losses. Income yields from investments may fluctuate and, in consequence, initial capital paid to make the investment may be used as part of that income yield. Some investments may not be readily realisable and it may be difficult to sell or realise those investments, similarly it may prove difficult for you to obtain reliable information about the value, or risks, to which such an investment is exposed.

Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information, opinions and estimates contained in this report reflect a judgement at its original date of publication by Grupo Santander and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. The value of securities and financial instruments is subject to exchange rate fluctuation that may have a positive or adverse effect on the price or income of such securities or financial instruments.

Information and opinions presented in this report have been obtained or derived from sources believed by Grupo Santander to be reliable, but Grupo Santander makes no representation as to their accuracy or completeness. Grupo Santander accepts no liability for loss arising from the use of the material presented in this report, except that this exclusion of liability does not apply to the extent that such liability arises under specific statutes or regulations applicable to Grupo Santander. This report is not to be relied upon in substitution for the exercise of independent judgment. Grupo Santander may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect the different assumptions, views and analytical methods of the analysts who prepared them and Grupo Santander is under no obligation to ensure that such other reports are brought to the attention of any recipient of this report.

See back cover of this report for further disclaimer disclosures.

Important Disclosures

The views expressed in this report accurately reflect the personal views of the Santander Investment Bolsa research team. In addition, the Santander Investment Bolsa research team has not and will not receive any compensation for providing a specific recommendation or view in this report: Colin Bennett and Miguel A. Gil

The analysts referenced in connection with the section for which he or she is responsible may have received or will receive compensation based upon, among other factors, the overall profitability of Grupo Santander, including profits derived from investment banking activities.

Grupo Santander may have a financial interest in the issuers/companies mentioned in this report, including a long or short position in their securities and/or options, futures or other derivative instruments based thereon or vice versa. Information relating to this is available upon request.

Grupo Santander may, in the last twelve months, have solicited or performed investment banking and/or underwriting activities and/or other services (including acting as adviser, manager, or lender) for any issuer/company referred to in this report which may have given rise to a payment or promise of a payment in relation to these services. Information relating to this is available upon request.

KEY TO INVESTMENT CODES*

		% of Companies	
Rating	**Definition**	**Covered with This Rating**	**Provided with Investment Banking Services in Past 12M**
Buy	Upside of more than 15%.	47.28	20.11
Hold	Upside of 10%-15%.	33.15	14.13
Underweight	Upside of less than 10%.	19.57	3.80
Under Review		0.00	0.00

NOTE: Given the recent volatility seen in the financial markets, the recommendation definitions are only indicative until further notice.
(*) Target prices set from January to June are for December 31 of the current year. Target prices set from July to December are for December 31 of the following year.

This report has been prepared by Santander Investment Bolsa, S.V., S.A. and is provided for information purposes only. Santander Investment Bolsa, S.V., S.A. is registered in Spain and is authorised and regulated by the CNMV.

This report is issued in the United States by Santander Investment Securities Inc. (“SIS”), in Spain by Santander Investment Bolsa, S.V., S.A. (“SIB”) under the supervision of the CNMV and in the United Kingdom by Banco Santander, S.A., London Branch (“Santander London”). SIS is registered in the United States and is a member of FINRA. Santander London is registered in the United Kingdom and subject to limited regulation by the Financial Services Authority, UK (“FSA”). SIS, SIB and Santander London are members of Grupo Santander. A list of authorised legal entities within Grupo Santander is available upon request.

This material constitutes “investment research” for the purposes of the Markets in Financial Instruments Directive and as such contains an objective or independent explanation of the matters contained in the material. Any recommendations contained in this document must not be relied upon as investment advice based on the recipient’s personal circumstances. The information and opinions contained in this report have been obtained from, or are based on, public sources believed to be reliable, but no representation or warranty, express or implied, is made that such information is accurate, complete or up to date and it should not be relied upon as such. Furthermore, this report does not constitute a prospectus or other offering document or an offer or solicitation to buy or sell any securities or other investment. Information and opinions contained in the report are published for the assistance of recipients, but are not to be relied upon as authoritative or taken in substitution for the exercise of judgement by any recipient, are subject to change without notice and not intended to provide the sole basis of any evaluation of the instruments discussed herein.

Any reference to past performance should not be taken as an indication of future performance. This report is for the use of intended recipients only and may not be reproduced (in whole or in part) or delivered or transmitted to any other person without the prior written consent of SIB.

Investors should seek financial advice regarding the appropriateness of investing in financial instruments and implementing investment strategies discussed or recommended in this report and should understand that statements regarding future prospects may not be realised. Any decision to purchase or subscribe for securities in any offering must be based solely on existing public information on such security or the information in the prospectus or other offering document issued in connection with such offering, and not on this report.

The material in this research report is general information intended for recipients who understand the risks associated with investment. It does not take into account whether an investment, course of action, or associated risks are suitable for the recipient. Furthermore, this document is intended to be used by market professionals (eligible counterparties and professional clients but not retail clients). Retail clients must not rely on this document.

To the fullest extent permitted by law, no Grupo Santander company accepts any liability whatsoever (including in negligence) for any direct or consequential loss arising from any use of or reliance on material contained in this report. All estimates and opinions included in this report are made as of the date of this report. Unless otherwise indicated in this report, there is no intention to update this report.

SIB may make a market in, or may, as principal or agent, buy or sell securities of the issuers mentioned in this report or derivatives thereon. SIB may have a financial interest in the issuers mentioned in this report, including a long or short position in their securities and/or options, futures or other derivative instruments based thereon, or vice versa.

Important Disclosures (continued from previous page)

SIB may receive or intend to seek compensation for investment banking services in the next three months from or in relation to an issuer mentioned in this report. Any issuer mentioned in this report may have been provided with sections of this report prior to its publication in order to verify its factual accuracy.

ADDITIONAL INFORMATION

SIB or any of its salespeople, traders and other professionals may provide oral or written market commentary or trading strategies to its clients that reflect opinions that are contrary to the opinions expressed herein. Also, in certain cases, a security may be covered by more than one analyst within Grupo Santander. Each of these analysts may use different methodologies to value the security and, as a result, the recommendations may differ and the price targets and estimates of each may vary widely. Furthermore, SIB's trading and investment businesses may make investment decisions that are inconsistent with the recommendations expressed herein.

No part of this report may be copied, conveyed, distributed or furnished to any person or entity in any country (or persons or entities in the same) in which its distribution is prohibited by law. Failure to comply with these restrictions may breach the laws of the relevant jurisdiction.

Investment research issued by SIB is prepared in accordance with Grupo Santander policies for managing conflicts of interest. In relation to the production of investment research, Grupo Santander has internal rules of conduct that contain, among other things, procedures to prevent conflicts of interest including Chinese Walls and, where appropriate, establishing specific restrictions on research activity. Information concerning the management of conflicts of interest and the internal rules of conduct are available on request.

COUNTRY- AND REGION-SPECIFIC DISCLOSURES

United Kingdom (UK): Unless specified to the contrary, this report has been distributed in the UK by Banco Santander, S.A. London Branch and has been prepared in accordance with Grupo Santander's policies for managing conflicts of interest arising as a result of publication and distribution of investment research. This report has been issued in the UK only to market professionals (professional or eligible counterparties but not retail clients) or persons of a kind described in Article 19 (5), 38, 47 and 49 of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (all such persons being referred to as "relevant persons"). The content must not be acted on or relied on by persons who are not relevant persons. Any investment or investment activity to which this communication relates is only available to relevant persons and will be engaged in only with relevant persons. **European Economic Area (EEA):** Unless specified to the contrary, this report has been approved for distribution in the EEA by SIB. Investment research issued by SIB has been prepared in accordance with Grupo Santander's policies for managing conflicts of interest arising as a result of publication and distribution of investment research. Many European regulators require that a firm establish, implement and maintain such a policy. Any investment or investment activity to which this document relates is only regarded as being provided to professional investors (or equivalent) in their home jurisdiction. **United States of America (US):** This report is being distributed to US persons by Santander Investment Securities Inc ("SIS") or by a subsidiary or affiliate of SIS that is not registered as a US broker dealer, to US major institutional investors only. Any US recipient of this report (other than a registered broker-dealer or a bank acting in a broker-dealer capacity) that would like to effect any transaction in any security or issuer discussed herein should contact and place orders in the US with the company distributing the research, SIS at (212) 692-2550, which, without in any way limiting the foregoing, accepts responsibility (solely for purposes of and within the meaning of Rule 15a-6 under the US Securities Exchange Act of 1934) under this report and its dissemination in the US. US recipients of this report should be advised that this research has been produced by a non-member affiliate of SIS and, therefore, by rule, not all disclosures required under NASD Rule 2711 apply. **Hong Kong (HK):** This report is being distributed in Hong Kong by a subsidiary or affiliate of Banco Santander, S.A. Hong Kong Branch, a branch of Banco Santander, S.A. whose head office is in Spain. The 1% ownership disclosure satisfies the requirements under Paragraph 16.5(a) of the Hong Kong Code of Conduct for persons licensed by or registered with the Securities and Futures Commission, HK. Banco Santander, S.A. Hong Kong Branch is regulated as a Registered Institution by the Hong Kong Monetary Authority for the conduct of Advising and Dealing in Securities (Regulated Activity Type 4 and 1 respectively) under the Securities and Futures Ordinance. The recipient of this material must not distribute it to any third party without the prior written consent of Banco Santander, S.A. **Japan (JP):** This report has been considered and distributed in Japan to Japanese-based investors by a subsidiary or affiliate of Banco Santander, S.A. - Tokyo Representative Office, not registered as a financial instruments firm in Japan, and to certain financial institutions defined by article 17-3, item 1 of the Financial Instruments and Exchange Law Enforcement Order. Some of the foreign securities stated in this report are not disclosed according to the Financial Instruments and Exchange Law of Japan. There is a risk that a loss may occur due to a change in the price of the shares in the case of share trading and that a loss may occur due to the exchange rate in the case of foreign share trading. **China (CH):** This report is being distributed in China by a subsidiary or affiliate of Banco Santander, S.A. Shanghai Branch ("Santander Shanghai"). Santander Shanghai or its affiliates may have a holding in any of the securities discussed in this report; for securities where the holding is greater than 1%, the specific holding is disclosed in the Important Disclosures section above. **Singapore (SI):** This report is issued and distributed in Singapore by Banco Santander, S.A. (Representative Office), authorised by the Monetary Authority of Singapore. It is not intended for distribution to any persons other than accredited investors, as defined in the Securities and Futures Act (Chapter 289 of Singapore), and persons whose business involves the acquisition or disposal of, or the holding of capital markets products (as defined in the Securities and Futures Act (Chapter 289 of Singapore)).
For further country and region specific disclosures please refer to Santander Investment Bolsa, S.V., S.A.

LOCAL OFFICES

Madrid Tel: 34-91-257-2309 Fax: 34-91-257-1811	**London** Tel: 44-207-756-6605 Fax: 44-207-756-5033	**Frankfurt** Tel: 49 6959 67-6403 Fax: 49 6959 67-6407	**Milan** Tel: 39-02-8542-09810 Fax: 39-02-8606-71648	**Lisbon** Tel: 351-21-381-6580 Fax: 351-21-389-3680
Brussels Tel: 32 2 286 5447 Fax: 32 2 230 6724	**Paris** Tel: 33 15353 7000 Fax: 33 15353 7060	**New York** Tel: 212-692-2550 Fax: 212-407-4540	**Hong Kong** Tel: 852 3921-3091 Fax: (852) 2101-2000	**Tokyo** Tel: 813-3211-0356 Fax: 813-3211-0362
Bogotá Tel: 571-644-8006 Fax: 571-592-0638	**Buenos Aires** Tel: 54114-341-1052 Fax: 54114-341-1226	**Caracas** Tel: 582-401-4306 Fax: 582-401-4219	**Lima** Tel: 511-215-8100 Fax: 511-215-8185	**Mexico City** Tel: 5255-5629-5040 Fax: 5255-5629-5846
Santiago Tel: 562-336-3300 Fax: 562-697-3869	**São Paulo** Tel: 5511-5538-8226 Fax: 5511-5538-8407			