

Volatility Investing Handbook


Topic Paper

Equity Derivatives Technical Study

Variance swaps and beyond

- The development of **new derivatives instruments** and an improvement in their liquidity means **volatility is becoming almost as easy to trade as stocks or bonds**.
- The goal of this research note is to explain **different ways of taking views on volatility**, starting with classic straddle/strangle positions and delta-hedging of options.
- Given the **weaknesses of these traditional volatility investing methods**, we introduce variance swaps, the first product to provide investors with pure volatility exposure.
- We show how **variance swaps** can be **replicated using a continuum of options**, both intuitively and mathematically. We provide an empirical description of variance strike prices and explain the presence of a variance risk premium.
- **We suggest several volatility-based strategies** ranging from the arbitrage of realized vs. implied volatilities and dispersion trading to the hedging of structured products or hedge fund strategies.
- We finally introduce a **new – or third – generation of volatility products** including **forward-start variance swaps, gamma swaps, corridor variance swaps and conditional variance swaps**. Thanks to their specific payoffs, these new products allow investors to bet on particular aspects of the volatility curve such as the smile or term structure and to trade dispersion efficiently.

Variance and gamma swap strike prices for the €Stoxx 50, Nikkei 225 and S&P 500



Source – BNP Paribas

Nicolas Mougeot
nicolas.mougeot@bnpparibas.com
+33 (0) 1 55 77 53 89

Table of contents

Executive summary	3
Volatility investing: an old story with some recent innovations	4
Volatility investing: the traditional way	5
Defining volatility	5
Taking vol positions by using straddles/strangles	5
Taking vol positions by delta-hedging options	6
The emergence of variance swaps and their valuation	8
The variance swap mechanism	8
Pricing: the intuitive approach	10
Pricing: Carr and Madan's formal approach	11
Why not volatility swaps?	13
Characteristics of volatility	14
The behavior of variance swap strike prices	15
Variance swap strike prices as a forecast of future realized volatility	16
The variance risk premium	17
Why use variance swaps?	19
Implied/realized volatility arbitrage	19
Volatility pairs and dispersion trading	20
Hedging structured products	22
Immunizing volatility risk of hedge fund strategies	23
Third-generation volatility products	24
Gamma swaps	24
Orderly dispersion trading	26
Forward-start variance swaps	26
Corridor variance swaps	27
Up and down corridor variance swaps	27
Up and down conditional variance swaps	29
Riding the smile	29
Correlation trading	31
Conclusion	34
Appendix	35
Hedging options when volatility is unknown	35
Valuing gamma swaps	36
The Derman <i>et al.</i> replication of variance swaps and its extension to gamma swaps	38
Example of capped variance-swap termsheet	40
References and further reading	42

Executive summary

The goal of this paper is to introduce investors to sophisticated derivatives instruments that allow them to take views on volatility. Despite the fact that volatility-sensitive derivatives have been traded for years if not centuries, it is only recently that products giving pure exposure to volatility have appeared.

We start our overview by discussing the weaknesses of “traditional” methods of volatility investing. Although buying/selling straddles is an easy way to bet on volatility, the delta – i.e. the portfolio’s sensitivity to stock returns – ceases to be null once the stock moves away from its initial value.

One way to circumvent this problem is to delta-hedge the positions. However, as we shall demonstrate below, delta-hedging options does not yield pure exposure to volatility since the trader/investor then faces a further vega risk as well as a model/path dependency risk.

Variance swaps provide a pure view on volatility since they pay the difference between future realized variance and a pre-defined strike price. The rapid development of variance swaps reflects the simplicity with which they can be valued: under certain non-restrictive assumptions, the variance swap strike price can be shown to be equal to the value of an options portfolio that uses a continuum of strike prices and is inversely weighted by the square of the options’ strike prices.

Variance swaps can be used in many ways, ranging from arbitraging realized vs. implied volatility and dispersion trading, to hedging structured products or hedge fund strategies.

We also describe other volatility products that have been developed in recent years. Gamma swaps are similar to variance swaps but with a notional that is a function of the asset price. They have several advantages over variance swaps since they do not require any caps and are a more efficient tool for dispersion trading. Finally, we present derivatives instruments such as conditional variance swaps or corridor variance swaps that allow investors to make asymmetric bets on volatility and take positions on the skew and the smile.

Volatility products and their various uses

	Straddle	Delta-hedged option	Variance swap	Gamma swap	Conditional variance swap	Corridor variance swap	Correlation swap
Volatility bet	+	+	++	++	++	++	-
Volatility hedging	-	-	++	+	+	++	-
Dispersion trading	+	-	++	++	-	-	-
Correlation trading	-	-	+	++	-	-	++
Asymetric vol bet	-	-	-	+	++	++	-
Smile trading	-	-		++	+	+	-

Source – Don M. Chance, BNP PARIBAS Arbitrage

Volatility investing: an old story with some recent innovations

Investing in volatility contracts or volatility-sensitive products is nothing new. As Don M. Chance points out, the use of options dates back to 1700 B.C., while the first derivatives market - the Royal Stock Exchange in London - started forward trading in 1637. However, the growth of derivatives instruments has long been linked to the parallel development of mathematical tools for efficient pricing and hedging. Despite the breakthroughs presented by Bachelier in his 1900 PhD thesis, trading in equity options only began after the Black & Scholes model was published in 1973.

Contracts based on variance, such as variance swaps were already being mentioned in the 1990s, but they only really took off after Carr and Madan developed an original pricing method in 2001.

While the emergence of variance swaps has allowed investors to take a pure position on volatility without taking other risks, a third generation of volatility products, including gamma corridor and conditional variance swaps, has appeared. These products provide investors with tools for taking positions on the skew/smile and efficiently trading dispersion.

The goal of this paper is therefore to explain the methods of trading volatility, from straddles to variance swaps and third-generation products.

Volatility products have historically emerged through the development of mathematical solutions for efficient pricing and hedging.

A short history of derivatives products since 1700 B.C.²

1700 B.C.	Genesis, Chapter 41. According to Joseph's advice, an Egyptian Pharaoh, anticipating seven years of feast followed by seven years of famine, executes hedge by storing corn. Joseph is put in charge of administering the program.
580 B.C.	Thales the Milesian (from Aristotle's <i>Politics</i>) buys options on olive presses and makes a killing off the bumper olive crop.
1570	The Royal Stock Exchange opens in London for forward contracting
1637	Dutch Tulip Bulb Mania; Defaults on Tulip bulb contingent forward contracts
1690	Options begin trading on securities in London
1900	Louis Bachelier derives the first option model in his doctoral dissertation. Commodity prices are modeled as arithmetic Brownian motions
1974	Black and Scholes publish their seminal article on option pricing while Merton publishes article on rational and mathematical of option pricing.
1975	Equity options begin trading at American and Philadelphia Stock Exchange
90ies	Several volatility contracts including variance swaps emerge.
2001	Carr and Madan make variance swaps pricing easier with their article "Towards a theory of Volatility Trading"
Nowadays	Issuance of 3 rd generation volatility products such as gamma swaps, corridor variance swaps and conditional variance swaps

Source – Don M. Chance, BNP PARIBAS Arbitrage

² See "A Chronology of Derivatives" by Don M. Chance for an extensive chronology from 1700 B.C. to 1995.

Volatility investing: the traditional way

Or how to take a position on volatility inefficiently

Defining volatility

First of all, we need to define how volatility is measured. Throughout this report, we define volatility as the annualized standard deviation – noted σ – of the log of the daily return of the stock (or index) price, and variance as the square of the standard-deviation³.

$$\sigma^2 = \frac{252}{T} \sum_{i=1}^T (\ln(S_i/S_{i-1}))^2 \quad \text{and} \quad \sigma = \sqrt{\frac{252}{T} \sum_{i=1}^T (\ln(S_i/S_{i-1}))^2}$$


Volatility is usually defined as the sum of squared asset returns.

Both are good measures of stock variability. However, as we shall see below, standard deviation is a more meaningful measure of volatility, given that it is measured in the same units as stock return. However, most of the volatility products we present here are priced in terms of variance, reflecting the fact that variance-related products are generally easier both to value and replicate (with the help of vanilla options). We also explain why they are also more useful to traders.

Taking vol positions by using straddles/strangles

The simplest way to obtain volatility exposure is to buy or sell straddles or strangles. A straddle involves buying an ATM call and an ATM put, while a strangle is a combination of one OTM call and one OTM put.

Straddle and strangle profiles



Source: BNP Paribas

Investors who buy straddles are taking a bet on the stock price moving up or down by a wide margin. Since straddles are composed of a call and a put, their

³ Note that standard deviation is calculated without the drift term

Straddles and strangles are the easiest way to invest in volatility...

...but certainly not the most efficient one.

value rises with volatility and as such, they provide a means for investors to take positions on future realized volatility and on changes in implied volatility. Since the two options are at-the-money, straddles are initially quite expensive. In view of this, one way to decrease the premium paid upfront is to buy strangles since they combine two cheaper out-of-the-money options. However, an investor buying a strangle would require the stock price to move more in order to make money.

However, straddles and strangles do not provide pure exposure to volatility. For example, let us assume that Stock A has an initial price of \$100 and that the investor buys a 3-month straddle. Even if the stock price moves sharply during the 3 months but ends up at \$100 at maturity, the straddle will be worthless at maturity. Of course, had the implied volatility also risen, buying the straddle would have been profitable if it had been sold before maturity.

Furthermore, once the stock price has moved away from its initial position (here \$100) the straddle delta is no longer null, as shown by the above graph. The straddle then becomes sensitive to price movements and the same holds true for the strangle.

Taking vol positions by delta-hedging options

A possible remedy to this weakness lies in selling the option and delta-hedging the position.

At time t , the investor:

- sells an option, call or put, denoted by $V(S_t, T; \sigma_i)$, with maturity T , strike K on the stock S and an initial implied volatility σ_i
- buys Δ stocks in order to delta-hedge the position
- borrows $-V(S_t, T; \sigma_i) + \Delta S$ to finance the position

One step beyond: delta-hedging options.

In order to eliminate the sensitivity to stock prices and obtain pure exposure to volatility risk, the amount of stocks held, Δ , is periodically re-adjusted to ensure the portfolio's sensitivity to the stock price remains null. At time t , Δ is thus equal to:

$$\Delta_t = \frac{\partial}{\partial S} V(S_t, 0; \sigma)$$

We assume that the option is delta-hedged at a constant implied volatility σ_h . The $P\&L$ resulting from delta-hedging an option is by definition equal to the final cost of the option minus its initial cost minus the cost of delta-hedging the position. We show in the appendix that this $P\&L$ can be broken up into three components:

$$P\&L = \underbrace{(\sigma_i^2 - \hat{\sigma}^2)Tg_0}_1 + \underbrace{(\hat{\sigma}^2 - \sigma_h^2)T \left[g_0 - \frac{1}{T} \int_0^T g_t dt \right]}_2 + \underbrace{\int_0^T (\hat{\sigma}^2 - \sigma_t^2) g_t dt}_3$$

where g_t is by definition equal to:

$$g_t = \frac{e^{r(T-t)}}{2\sigma_h} \frac{\partial}{\partial \sigma} V(S_t, T-t; \sigma_h)$$

According to Blanc⁴, these three components may be described as:

I: A “Variance risk” component or a variance swap exposure for a notional amount equal to Tg_0 and a variance strike equal to the square of the option implied volatility.

II : A “Vega risk” factor, which stems from the fact that the option is hedged at the implied volatility σ_h instead of at the realized volatility $\hat{\sigma}$. This term is indeed null if the trader is able to hedge at the realized – but unknown – volatility.

III: A “Volatility path dependency risk” or “model risk” factor that depends on the historical behavior of realized volatility. Under the Black & Scholes assumption, the instantaneous volatility σ_t is constant and thus equal to the realized volatility between time t and T . However, should the volatility vary over time, $(\hat{\sigma}^2 - \sigma_t^2)$ will no longer be zero. As g_t is a decreasing function of time to maturity, this term will be positive if instantaneous – or intraday – volatility rises during the life of the option. The term also depends on the true distribution of stock returns.

Based on realistic simulations, Blanc also shows that variance risk only represents 52% of the total P&L resulting from delta-hedging an option. As a result, delta-hedging options does not provide pure exposure to volatility, given that the P&L not only depends on variance risk but also on a vega risk which itself results from the fact that risk cannot be hedged at the - unknown future - realized volatility and that volatility may not be constant over time.

Indeed, delta-hedging options yields further risk sources not indicated above. The above analysis was done without taking into account dividends and by assuming a constant interest rate. In practice, traders face the risk of unknown dividends being paid during the life of the option, while with interest rates liable to vary over time, the option vega may change if the interest rate changes.

Furthermore, options delta-hedging is impacted by transaction costs and liquidity issues. The above analysis does not take into account the fact that trading stocks is costly and that certain stocks and indices may lack liquidity.

These issues paved the way for the introduction of new derivatives instruments that enable investors to take a view on volatility without bearing any other risks.

Options delta-hedging implies three different sources of risk:

- **Variance risk**
- **Vega risk**
- **Volatility path dependency risk or model risk.**

⁴ For a discrete time version of the demonstration, see N. Blanc, Index Variance Arbitrage : Arbitrating Component Correlation, BNP Paribas technical studies, 2004

The emergence of variance swaps and their valuation

Or how variance swaps help in obtaining pure volatility exposure

Variance swaps emerged in the 1990s as a means of circumventing the issues raised by taking volatility positions through the purchase of straddles/strangles or through delta-hedging options. Although variance swaps were initially considered as exotic products, over recent years, they have become flow products as a result of certain features that we will describe in greater detail in this section. These are:

- simple payoffs
- simple replication via portfolios of vanilla options.


The variance swap mechanism

Variance swaps are forward contracts that pay at maturity the difference between the realized variance of an underlying asset (stock, index...) and the initially defined variance strike price K_{Var} , times a notional N .

Cash-flow of a variance swap at maturity

A simple pay-off:

$$Pay-off = (\sigma^2 - K_{Var}) * N$$



Source – BNP Paribas

Although there is no formally defined market convention for calculating the variance of an asset, variance is usually calculated as:

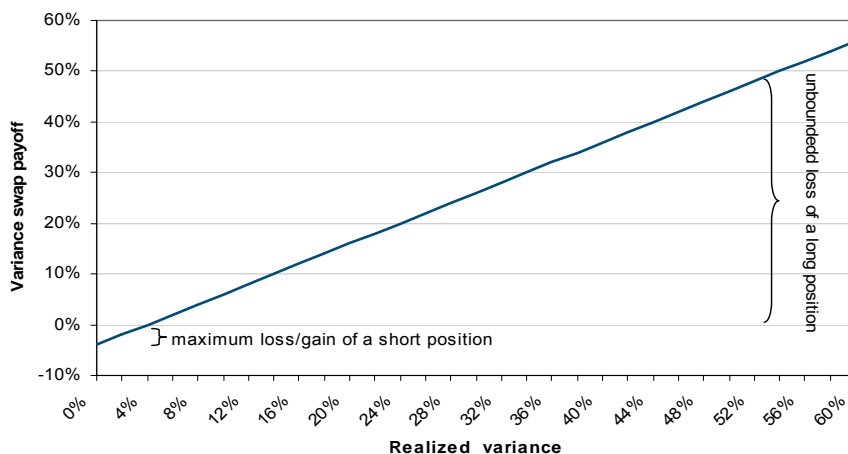
$$\sigma^2 = \frac{252}{T} \sum_{i=1}^T (\ln(S_i/S_{i-1}))^2$$

Where S_i is the closing price of the asset, T the number of days in the observation period and \ln the natural logarithm. Also, variance swaps are usually quoted in terms of squared volatility, e.g. $(20\%)^2$ since volatility is more economically meaningful than variance.

Note that there is no correction for dividend payment in the above formula. Should a stock pay a dividend, the final payout will not be adjusted for the jump implied by the dividend. However, variance swaps can be structured in such a way as to adjust the above formula for dividends. One should also bear in mind that variance swaps are not exempt from counterparty risks. Since the product is OTC, both counterparties face the risk of the other going bankrupt.

Another feature of variance swaps is that taking short positions on a variance swap on a single stock may be risky. Whereas the total loss on short positions is unbounded, the maximum loss on long positions is equal to the notional times the variance strike.

Variance swap payoff with $K_{VAR}=(20\%)^2$



Source – BNP Paribas

Should the stock move sharply because of a crash, an M&A announcement or any other event, its volatility may spike. This is important as there is no limit to the potential loss of being short on a variance swap. As a result, variance swaps are often capped in such a way that:

Variance swaps may be capped in order to limit the maximum loss faced by the trader/investor.

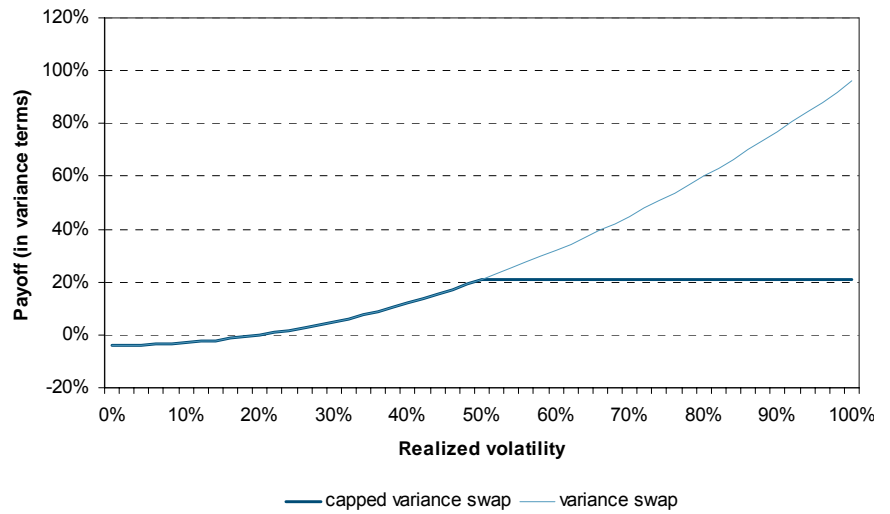
Variance swap buyer pays:

$$\text{Notional} * \text{Max} [0, K_{VAR} - \text{Min}(m^2 * K_{VAR}; (\text{Realized Volatility})^2)]$$

Variance swap seller pays:

$$\text{Notional} * \text{Max}[0, \text{Min}(m^2 * K_{VAR}; (\text{Realized Volatility})^2) - K_{VAR}]$$

m is the multiplier setting the cap. The graph below compares the payoffs from capped and uncapped variance swaps. The maximum gain (and thus loss for the counterparty) is limited and thereby facilitates the hedging of the instrument. As we will demonstrate in the next section, a variance swap can in theory be replicated by a portfolio of options with a continuum of strike prices. In practice, however, there is no liquidity for options whose strike prices are far from ATM. Traders are therefore unable to completely hedge variance swaps, whereas capped variance swaps do not require a continuum of strike prices.

Comparison of capped and uncapped variance swaps, $K_{VAR}=20\%$, $m=2.5$


Source – BNP Paribas

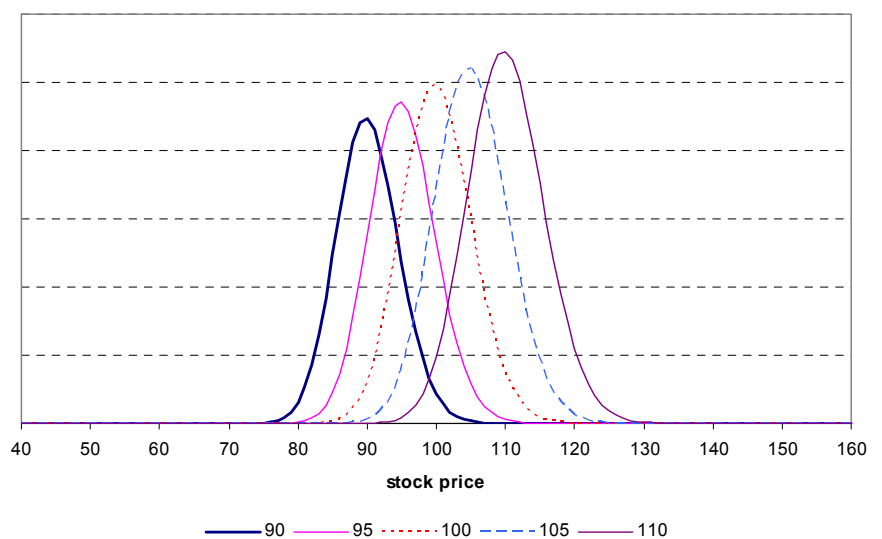
Pricing: the intuitive approach

Before examining variance swap pricing in detail, we shall explain the logic behind the mathematical derivation.

Firstly, variance swap pay-off is by definition a function of the variance and is independent of the stock price level. However, the sensitivity of an option to variance depends on the stock price level. The sensitivity of an option to the variance, or variance vega, is centered around the strike price and will thus change daily according to changes in the stock price level.

As the graph below shows, the variance vega declines as the stock price moves away from the strike price and is also an increasing function of the strike price. The goal is therefore to create an options portfolio with a constant variance vega.

The strike price may be determined by a static portfolio of options inversely weighted by the square of their strike price.

Variance vega of 3-month options with different strike prices


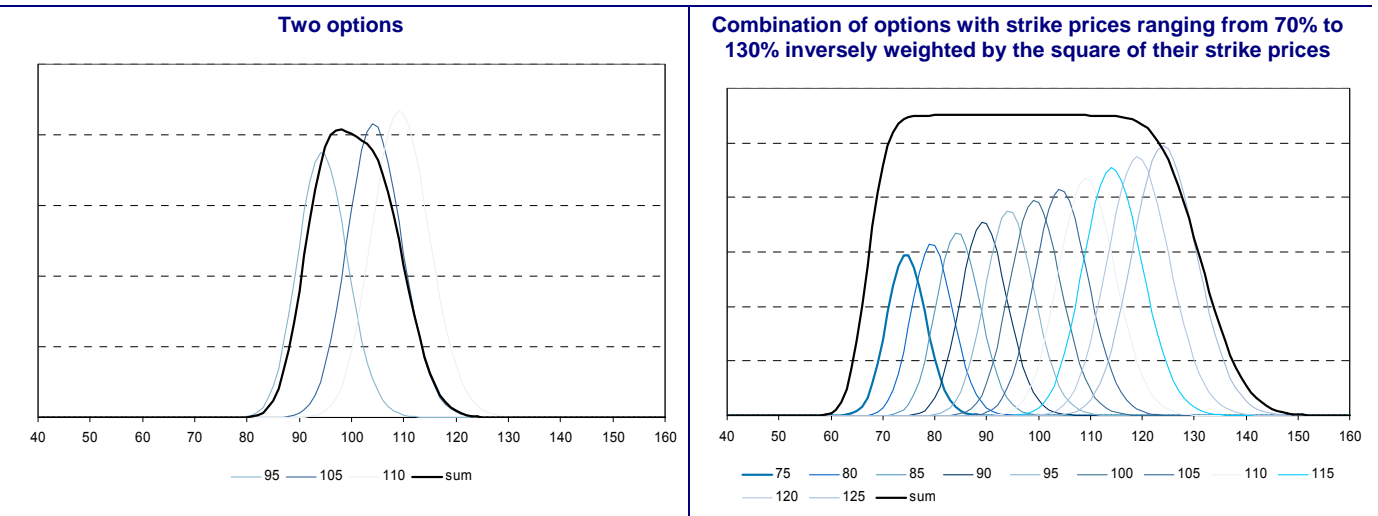
Source – BNP Paribas

This can be done by investing in a portfolio of options inversely weighted by the square of their strike prices. The following two graphs display the variance vega of:

- 1- a portfolio composed of 2 options with strike prices of 95% and 105%, and weighted $1/(95)^2$ and $1/(105)^2$ respectively, and
- 2- a portfolio composed of 13 options with strike prices ranging from 70 to 130% and weighted from $1/(70)^2$ to $1/(130)^2$.

The addition of options with strike prices away from ATM flattens variance vega and thus makes the portfolio sensitive to variance but invariant to stock price.

Option-portfolio variance vega



Source: BNP Paribas

Carr and Madan suggest an interesting *model-free* method for valuing variance swaps.

Pricing: Carr and Madan's formal approach

Carr and Madan suggest a formal method for pricing variance swaps that has the advantage of requiring very few assumptions about stock price dynamics. Instead of defining a process for the stock price, Carr and Madan only assume that markets are complete and that trading can take place continuously. As a result, their results hold for restrictive assumptions such as those underlying the Black & Scholes model and can also be extended to other models such as stochastic volatility. Although this section is rather mathematical it describes one of the greatest results in derivatives research and one that has largely contributed to the growth of the variance swaps market.

Firstly, Ito's Lemma states that any smooth function $f(F_t)$ can be rewritten as:

$$f(F_T) = f(F_0) + \int_0^T f'(F_t) dF_t + \frac{1}{2} \int_0^T F_t^2 f''(F_t) \sigma_t^2 dt \quad (1)$$

Let us consider the following function:

$$f(F_t) = \ln(F_0/F_t) + \frac{F_t}{F_0} - 1$$

Where F_t stands for the futures price. This function has a slope and a value equal to zero when $F_t = F_0$.

Therefore, we have:

$$\frac{1}{2} \int_0^T \sigma_t^2 dt = \ln(F_0/F_T) + \frac{F_T}{F_0} - 1 - \int_0^T \left(\frac{1}{F_0} - \frac{1}{F_t} \right) dF_t \quad (2)$$

Carr and Madan further assume that a market exists for futures options of all strikes. In this case, they show that any payoff $f(F_T)$ of the futures price F_T can be broken up into:

$$\begin{aligned} f(F_T) = & \underbrace{f(\kappa)}_1 + \underbrace{f'(\kappa)[(F_T - \kappa)^+ - (\kappa - F_T)^+]}_2 \\ & + \underbrace{\int_0^\kappa f''(\kappa)(\kappa - F_T)^+ d\kappa}_3 + \underbrace{\int_\kappa^\infty f''(\kappa)(F_T - \kappa)^+ d\kappa}_4 \end{aligned} \quad (3)$$

Where κ is an arbitrary number. As Carr and Madan point out, the above terms can be interpreted as:

- 1- a position in $f(\kappa)$ discount bond
- 2- a position in $f'(\kappa)$ calls with a strike price equal to κ minus $f'(\kappa)$ puts with the same strike price
- 3- a static position in $f''(\kappa)dK$ puts at all strikes less than κ
- 4- a static position in $f''(\kappa)dK$ calls at all strikes greater than κ

In the absence of arbitrage, the above breakdown must prevail for initial values. Therefore, the initial value of the payoff is equal to:

$$V_f^0 = f(\kappa)B_0 + f'(\kappa)[C_0(\kappa) - P_0(\kappa)] + \int_0^\kappa f''(\kappa)P_0(\kappa)dK + \int_\kappa^\infty f''(\kappa)C_0(\kappa)dK$$

Carr and Madan thus prove that an arbitrary pay-off can be obtained from bond and option prices without making strong assumptions about the stochastic process driving the stock price⁵.

Applying equation (3) to the function $f(F_t) = \ln(F_0/F_t) + (F_t/F_0) - 1$ and setting $\kappa = F_0$, we get:

$$\ln(F_0/F_T) + \frac{F_T}{F_0} - 1 = \int_0^{F_0} \frac{1}{K^2} (K - F_T)^+ dK + \int_{F_0}^\infty \frac{1}{K^2} (F_T - K)^+ dK$$

Therefore, in order to receive $\int_0^T \sigma_t^2 dt$ at time T , a trader should buy a continuum

of puts with strike prices ranging from 0 to F_0 and calls with strikes prices ranging from F_0 to infinity. The initial cost is equal to:

$$\int_0^{F_0} \frac{2}{K^2} P_0(K) dK + \int_{F_0}^\infty \frac{2}{K^2} C_0(K) dK \quad (A)$$

The trader further needs to roll a futures position, holding at t :

$$-2e^{-rt} \left(\frac{1}{F_0} - \frac{1}{F_t} \right) \quad (B)$$

⁵ More complete proof of the result can be found in Peter Carr and Dilip Madan, 1998, 'Towards a Theory of Volatility Trading', in *Volatility: New Estimation Techniques for Pricing Derivatives*, pp. 417-427.

The net payoff of (A) and (B) at maturity T is:

$$\begin{aligned} & \int_0^{F_0} \frac{2}{K^2} (K - F_T)^+ dK + \int_{F_0}^{\infty} \frac{2}{K^2} (F_T - K)^+ dK - 2 \int_0^T \left(\frac{1}{F_0} - \frac{1}{F_t} \right) dF_t \\ &= 2 \left(\ln(F_0/F_T) + \frac{F_T}{F_0} - 1 \right) - 2 \int_0^T \left(\frac{1}{F_0} - \frac{1}{F_t} \right) dF_t \\ &= \int_0^T \sigma_t^2 dt \end{aligned}$$

Since the initial cost of achieving this strategy is given by (A), the fair forward value of the variance at time 0 should be equal to:

$$V_f[\sigma_{0,T}^2] = \frac{2e^{rT}}{T} \left[\int_0^{F_0} \frac{1}{K^2} P_0(K) dK + \int_{F_0}^{\infty} \frac{1}{K^2} C_0(K) dK \right]$$

Variance-swap strike prices may thus be replicated by a continuum of puts and calls inversely weighted by the square of their strike price. Note that the above valuation is model-free since we did not have to state a specific process for the dynamic of the stock price in order to derive the formula. It is also worth mentioning that the above formula provides a market-based estimator of future realized volatility.

Why not volatility swaps?

Volatility, not variance, is the most commonly used risk measure. First, it is measured in the same unit as stock returns. If returns are normally distributed, they should fall within ± 2 standard deviations or volatility away from the mean in 95% of cases. However, variance is less meaningful despite it being the square of the standard-deviation.

Nevertheless, variance swaps are much more widely traded than volatility swaps for two reasons: they are much easier to replicate and the P&L of a delta-hedged option is a linear function of variance and not of volatility. As shown above, variance swaps can be replicated with the help of a linear combination of options and a dynamic position in futures. Replicating volatility is more complex and requires a *non-linear* combination of derivative instruments.

Furthermore, let us recall that the P&L of a delta-hedged option is given by:

$$P \& L = (\sigma_i^2 - \hat{\sigma}^2) T g_0 + (\hat{\sigma}^2 - \sigma_i^2) T \left[g_0 - \frac{1}{T} \int_0^T g_t dt \right] + \int_0^T (\hat{\sigma}^2 - \sigma_t^2) g_t dt$$

Options traders are therefore linearly exposed to variance and are thus likely to be more interested in variance swaps. The same reasoning holds true for an investor wishing to hedge the volatility exposure of his options portfolio.

However, the strike price of a variance swap is often defined in terms of volatility as it is more economically meaningful. For example, a variance swap contract will specify that the volatility strike is equal to 20% which means that the actual variance strike price is equal to $(20\%)^2$.

Volatility swaps are more difficult to value and less useful for option hedging.

The variance exposure of a variance swap is equal to $\sigma^2 - K_{VAR}$ and its exposure to volatility may be expressed as:

$$\sigma - K_{VOL} \approx \frac{1}{2K_{VOL}}(\sigma^2 - K_{VAR})$$

Where K_{VOL} is the strike price measured in terms of volatility (i.e. the square root of the variance strike price).

Characteristics of volatility


The specific dynamics of volatility are of particular interest and these may be summarized in the following four points:

- it jumps when markets crash
- it tends to revert back towards its long-term mean
- it experiences high and low regimes
- it is usually negatively correlated with the underlying asset return

Volatility tends to revert back towards its long-term mean and jump when markets crash.

The following graphs display 6-week historical volatility for the VIX⁶ and the S&P 500. The graph on the left shows that volatility may move from around 20% to north of 100% when the market experiences a serious downturn such as in October 1987. Furthermore, as opposed to stock returns, volatility tends to revert back towards its mean and usually remains within a high or a low regime for a long period of time. S&P 500 volatility, for example, was in a low regime between 1992 and 1996 and then moved into a high regime until 2003.

VIX and S&P 500 6-week historical volatility




Source – BNP Paribas

The other empirical characteristic of volatility is that it is usually negatively correlated with the underlying asset return. The following graph plots the 6-week correlation between changes in the VIX index and S&P 500 daily returns. The correlation has been negative 98% of the time since January 1986. This

⁶ The Chicago Board Options Exchange SPX Volatility, or VIX, Index reflects a market estimate of future volatility, based on the weighted average of the implied volatilities for a wide range of strikes. 1st & 2nd month expirations are used until 8 days from expiration, then the 2nd and 3rd are used.

contradicts the Black & Scholes assumption of lognormal prices and highlights the fact that the distribution of stock returns is skewed.

6-week correlation between changes in VIX and S&P 500 daily returns



Source – BNP Paribas

Volatility is usually negatively correlated with equity returns.

The behavior of variance swap strike prices

The following table shows a few statistics based on the €-Stoxx 50 and calculated from August 2001 to August 2005. The 6-month variance swap strike is calculated using the method suggested by Derman et al. (see appendix for details), using option prices with strikes ranging from 40% to 160%.

First, let us compare the 6-month variance swap strike price with 6-month ATM implied volatility for European options on the €-Stoxx 50 index. Variance swaps would have traded an average of 1.11 volatility points above ATM implied volatility between August 2001 and August 2005, a difference that appears to be statistically significant.

Statistics on a 6-month variance swap on the €-Stoxx 50 from August 2001 to August 2005

	Var swap	ATM implied vol	Skew (90% - 100%)	Var swap - ATM implied vol	Var swap - realized vol
minimum	13.17%	12.43%	1.75%	0.31%	-19.86%
maximum	40.04%	39.52%	3.32%	2.74%	16.39%
average	26.28%	25.17%	2.40%	1.11%	1.11%


Source – BNP Paribas

Variance swap strike prices are higher than realized volatility on average.


We further describe the payoff from a variance swap, calculated as the difference between the swap's initial strike price (in volatility terms) and the volatility realized during the relevant period. The positive average comes from the existence of a variance risk premium, which we highlight in the next section.

Empirical analysis of a 6-month variance swap on the €Stoxx 50

Realized volatility, payoff and strike of a 6-month variance swap on the €Stoxx 50



ATM volatility vs. variance swap strike (in vol terms) for the €Stoxx 50



Source – BNP Paribas

Variance swap strike prices as a forecast of future realized volatility

By definition, the variance swap strike price should be a good proxy of future realized volatility. If this assertion holds true, regressing the variance swap strike price against realized volatility should yield a slope and a constant equal to one and zero, respectively.

The following table plots the regression of the €Stoxx 50's 1-year historical standard-deviation against its 1-year variance strike price (expressed in terms of volatility) and against 1-year ATM implied volatility. Since regression β is far from one in the case of variance swap strike prices, it does not provide an unbiased proxy for future realized volatility. The same conclusion holds for ATM implied volatility despite the fact that the constant and the slope are closer to zero and one, respectively.

Regression of €Stoxx 50 1-year realized volatility against its corresponding var swap strike price and ATM implied volatility⁷

	constant	T-stat (constant)	β	T-stat (β)	R-square
var swap	0.10	5.27	0.59	8.56	8%
ATM implied volatility	0.06	3.75	0.76	11.85	13%

Source – BNP Paribas

Variance swap strike prices are not a good proxy for future realized volatility.

One of the reasons why variance-swap strike prices do not provide a fair proxy of future realized volatility is again the presence of a variance risk premium.

⁷ β is the sensitivity of the variance-swap strike price (or ATM implied volatility) to realized volatility. The T-stat measures whether this parameter shows a statistically significant difference from zero: it needs to be below -1.96 or above 1.96 to be significant at the 95% confidence level. The R-square measures the overall explanatory power of the equation. An R-square of 40% means that realized volatility explains 40% of the variability of ATM implied volatility.

The variance risk premium

As already mentioned, sellers of variance swaps bear a bigger risk than buyers as their potential loss is unlimited. They should therefore be rewarded with a variance risk premium that is reflected in a variance strike price higher than realized variance on average⁸.

To check whether a variance risk premium exists, we follow Carr and Wu's suggestion: in the absence of a variance risk premium, the average variance swap strike price should be equal to the realized variance. The log ratio of the two - $\ln(\text{Var Swap}/\text{realized variance})$ - should thus be equal to zero.

The table below reports the t-statistics for $\ln(\text{var swap}/\text{realized variance})$ for the €-Stoxx 50 and its 10 largest constituents at present using data from March 2001 to September 2005. The t-stat is significant and positive both for the €-Stoxx 50 and for 9 out of its top 10 constituents, a situation denoting the presence of a significant variance risk premium at both the index and single stocks levels.

Variance risk is priced into variance swaps.


Testing for variance risk premium	
t-stat for the €Stoxx 50 and its 10 biggest constituents	
underlying asset	t-stat
€-Stoxx 50	8.70
Total SA	11.19
Sanofi-Aventis	6.62
Banco Santander Central Hispano SA	3.55
ENI SpA	8.84
Telefonica SA	59.23
Nokia OYJ	(4.03)
E.ON AG	2.53
Siemens AG	5.80
BNP Paribas	2.64
Banco Bilbao Vizcaya Argentaria SA	2.15

Source – BNP Paribas

Indeed, in a risk-neutral world, this variance risk premium should be equal to zero. A risk-neutral investor should sell a variance swap at his expected realized variance. In practice, however, the presence of risk aversion makes the variance risk premium positive.

Furthermore, the skew (calculated here as the 90%-100% implied volatility spread), indirectly measures the market's risk aversion. As a result, one way to test the effect of risk aversion on the variance risk premium is to check whether the difference between the variance-swap strike price and realized variance is a function of the skew. As shown by the following table, the higher the skew, the higher the variance risk premium.

⁸ Carr and Wu (2005) provide a thorough analysis of variance risk premium for several US stocks and indices. See also Driessen et al. (2005) for a comparison of variance risk premia for single stocks and indices.

Regression of €Stoxx 50 1-year skew (90% - 100%) against the spread between the 1-year var swap strike price and 1-year realized volatility


	constant	T-stat (constant)	β	T-stat (β)	R-square
var swap - realized vol	-0.30	-15.11	13.03	15.81	21.67%

Source – BNP Paribas

Why use variance swaps?

In this section, we show how variance swaps can be used to take positions on volatility or hedging it.

Implied/realized volatility arbitrage

The most obvious use of variance swaps is to bet on the difference between current implied and future realized volatilities. With variance swap strike prices defined by a combination of puts and calls, they may be viewed as a weighted average of implied volatility across strikes.

Variance swaps can thus be used for different strategies such as:

■ Implied vs. realized volatility arbitrage

If one expects future realized volatility to be above implied volatility, as measured by the strike price of a variance swap, investors can use a long position in a variance swap to take this view by means of a buy-and-hold strategy, as opposed to option delta-hedging which would require daily monitoring of the delta position.

■ Implied volatility term-structure arbitrage

One way to play an expected rise in volatility term structure is to enter into two different variance swaps with two different maturities. Let us assume for example that the variance strike price for a 1-year variance swap is currently equal to $(20\%)^2$, but that one expects the 6-month variance swap to be priced at $(30\%)^2$ in 6 months with the realized volatility remaining constant over the next 12 months. An investor can then take advantage of this expected steepening of volatility term structure by:

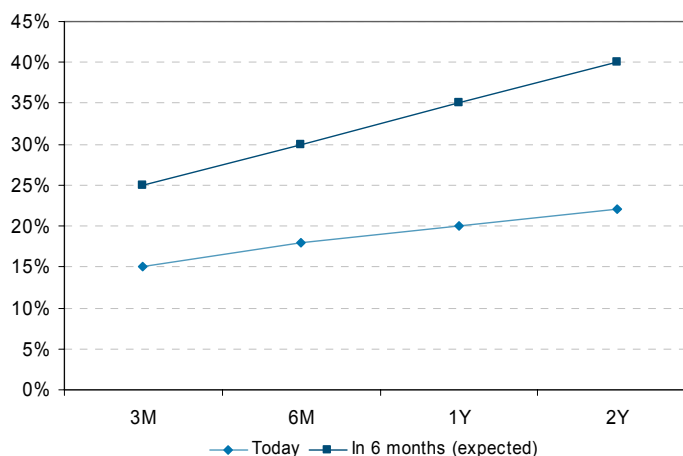
- buying a 1-year variance swap for a notional N today with a strike price equal to $K_{\text{var},0}=(20\%)^2$
- selling a 6-month variance swap for the same notional in 6 months with a strike price equal to $K_{\text{var},0}=(30\%)^2$

If the investor is right, the payoff will be equal to $N*(10\%)^2$. This type of strategy may suitably be executed via the use of forward-start variance swaps⁹ which we describe in more detail in the next section.

Variance swaps are an efficient tool for investing in volatility and arbitraging discrepancies between implied and realized volatilities.

⁹ The next section looks into forward variance swaps in greater detail.

Volatility term structure in terms of variance-swap strike price



Source – BNP Paribas

Variance swaps can be set up on indices as well as stocks, thereby allowing for easy implementation of volatility pairs trading.

Volatility pairs and dispersion trading


Here, we present two different single-stock volatility strategies, namely volatility pairs and dispersion trading.

■ Volatility pairs trading

Since the volatilities of two stocks within the same sector are usually driven by the same factors, their spread is often mean-reverting: should one implied volatility diverge from the other, it is likely to revert back. The following graph shows BMW and DaimlerChrysler's 3-month ATM implied volatilities in the past two years. Their levels are closely related and the implied volatility spread is reverting towards a mean around zero. As a result, the rise in DaimlerChrysler's implied volatility in May 2005 could have provided a trading opportunity that could have been exploited by using variance swaps.

Our implied volatility valuation model can also provide some ideas for implied volatility pairs trading. This model helps calculate fair values for the implied volatility of single stocks according to their beta, 5-year CDS, size and stock returns¹⁰.

¹⁰ For more details on this model, see "Predicting uncertainty: A new method for valuing implied volatility", BNP Paribas Technical Studies, published on October 14, 2004.

DaimlerChrysler and BMW's 6-month ATM implied volatility


Source – BNP Paribas

■ Dispersion trading

Dispersion trading consists of buying the volatility of an index and selling the volatility of its constituents according to their index weights. It is defined as¹¹:

$$Dispersion_t = \sum_{i=1}^N \alpha_{i,t} \sigma_{i,t}^2 - \sigma_{I,t}^2$$

Where $\alpha_{i,t}$ is the weight of stock i in the index. By definition, it changes over time as stock prices change and is equal to:

$$\alpha_{i,t} = n_i S_{i,t} / I_t$$

It thus depends on:


- the volatility of single stocks
- the correlation

As correlation tends to revert back towards its long-term mean, so does dispersion. Dispersion trades may be executed by buying a variance swap on the index and selling variance swaps on the constituents according to their weight in the index.

Nowadays, dispersion is usually quoted by mentioning the level of volatility of the index, the average weighted spread of volatility with its constituents and the level of correlation.

¹¹ Dispersion can also be expressed as the square root of the mentioned formula

€Stoxx 50 implied and realized dispersion



Source – BNP Paribas

Variance swaps can be useful tools for hedging the volatility risk of complex structured products such as CPPIs.

Hedging structured products

Given that structured products are by definition sensitive to volatility, variance swaps are natural tools for hedging volatility exposure. To illustrate this point, we take a look at CPPI or Constant Proportion Portfolio Insurance. A CPPI is a dynamic strategy that can be applied to stocks or other assets. The following paragraphs describe how CPPI works.

Let us suppose that the initial investment is 100 with a 5-year time horizon. The investor first determines a floor that is equal to the value of a 5-year bond. The cushion C_0 is equal to the initial investment I_0 minus the bond floor B_0 . The investor could invest C_0 in risky assets while investing the floor in a risk-free 5-year bond. In that case, the final value of his investment would be guaranteed against a decline below its initial value. However, the CPPI allows the investor to invest more on risky assets by an amount $m \cdot C_0$ where m stands for the multiplier (and is greater than 1). The higher the multiplier, the greater the participation in an appreciation of the risky asset.

The investment in the risky asset is dynamically adjusted in order to ensure a return of 100 at maturity: if the risky asset declines, exposure to the risky asset is consequently reduced. Once the total value of the portfolio reaches the bond value, the CPPI is said to have been cashed-out and all the money is invested in the risky bond as there is no more cushion.

CPPIs may be demonstrated to be negatively related to the volatility of the underlying asset for the following reasons:

- A high volatility increases the probability of cashing out
- A high volatility is often paired with a negative return.

On the other hand, option-based products such as ODBs, are usually positively related to volatility. On this basis, variance swaps could help structured-product managers to hedge their volatility risks efficiently.

Immunizing volatility risk of hedge fund strategies

Variance swaps are currently used by hedge fund managers to bet on volatility. They can also be used to protect against volatility risk. While numerous studies demonstrate the correlation between hedge fund strategies and equity market returns, few have focused on the effect of volatility.

The performance of hedge fund strategies are often affected by changes in volatility.

The following table shows that the correlation between the various hedge fund strategies and S&P 500 and €-Stoxx 50 realized volatility is usually far from being zero.

Convertible arbitrage, for example, benefits from a rise in volatility as higher volatility usually creates more arbitrage opportunities. Conversely, when volatility declines - as it has in the last two years - convertible hedge funds can improve their performance by selling variance swaps.

Conversely, hedge funds with strategies negatively correlated with volatility, such as event-driven arbitrage or distressed companies, could yield higher returns in a high volatility market if they buy variance swaps.

Correlation between CSFB/Tremont hedge fund indices and S&P 500 and €Stoxx 50 return and volatility

	CSFB/Tremont Hedge Fund Index	Convertible Arbitrage	Dedicated Short	Emerging Markets	Event Driven	Distressed	Global Macro	Long/Short Equity	Multi- Strategy
S&P 500 volatility	-31%	21%	21%	-29%	-45%	-41%	-9%	-32%	-4%
€Stoxx 50 volatility	-35%	13%	17%	-24%	-47%	-41%	-3%	-26%	-6%
S&P 500 return	41%	7%	-75%	63%	53%	44%	11%	44%	30%
€Stoxx 50 return	45%	4%	-63%	57%	49%	36%	17%	49%	28%

Source: BNP Paribas, CSFB/Tremont

Third-generation volatility products

Or how new products can further improve volatility trading by enabling asymmetric bets/hedges to be taken on volatility

Gamma swaps

As discussed previously, the gamma exposure of variance swaps is insensitive to the level of the underlying asset. In the event the stock price rises or declines, the gamma exposure depends solely on the initial value of the portfolio. Variance swaps are thus said to have constant “cash” gamma exposure.

However, it is often more useful to have a constant “share” gamma exposure than a “cash” one. In general, investors focus on the number of portfolio units they manage and not on the initial cash value of their portfolio. Gamma swaps are by definition products that answer to this need.

Gamma swaps have the following pay-off:

$$\text{Pay-off}_{\text{Gamma swap}} = (\text{Gamma} - K_{\text{Gamma}}) * N$$

Where $\text{Gamma} = \frac{252}{T} \sum_{i=1}^T \left[(\ln(S_i/S_{i-1}))^2 \frac{S_i}{S_0} \right]$, K_{Gamma} is the strike and N the

notional amount. In continuous-time, the gamma swap payoff is equal to:



$$\Gamma_{0,T} = \frac{1}{T} \int_0^T \sigma_t^2 S_t / S_0 dt$$

Gamma swaps are thus equivalent to variance swaps whose nominal is proportional to the level of the underlying asset.

■ The intuitive and formal approaches to gamma swap pricing

Pricing gamma swaps is as easy as pricing variance swaps. We have already shown that whereas the variance swap vega should be independent of the stock price, the gamma-swap vega should be a linear function of the stock price. The aim is to create an options portfolio whose vega will be a linear function of the stock price. As the following graph illustrates, an option portfolio using a continuum of strike prices and inversely weighted by their strike prices provides a vega that is linear with respect to the stock price.

Gamma swaps have a constant “share” gamma exposure which makes them more efficient tools for options hedging.

Option-portfolio variance vega
Two options

Combination of options with strike prices ranging from 70% to 130% (5% interval) and inversely weighted by their strike prices


Source: BNP Paribas

Carr and Madan's formal approach may also be applied to gamma-swap pricing. As shown in the appendix, the value – or the strike price – of a gamma swap is given by:


$$V_f[\Gamma_{0,T}] = \frac{2e^{2rT}}{TS_0} \left[\int_0^{F_0} \frac{1}{K} P_0(K) dK + \int_{F_0}^{\infty} \frac{1}{K} C_0(K) dK \right]$$

The gamma-swap's strike price may thus be replicated by a continuum of puts and calls inversely weighted by their strike prices.

Gamma-swap strike prices may be replicated by an options portfolio with a continuum of strike prices inversely weighted by their strike prices.

Empirical differences between gamma swaps and variance swaps

By virtue of their payoff and insofar as squared returns $(\ln(S_i/S_{i-1}))^2$ are weighted by the performance of the stock S_i/S_0 , gamma swaps underweight big downward index move relative to variance swaps. This means that if the distribution of stock returns is skewed to the left, gamma swaps minimize the effect of a crash, thereby making it easier for the trader to hedge. In this case, hedging does not require additional caps, unlike variance swaps which need to be capped.

6-month gamma and variance swaps' strike prices on the €Stoxx 50 index


Source: BNP Paribas

Gamma swap strike prices should thus be lower than variance swap strike prices. The following graphs show the strike prices of 6-month variance and gamma swaps on the €-Stoxx 50. The gamma swap strike price would systematically have been slightly lower than the variance swap strike price by 1.03 volatility points on average.

The gamma swap payoff is slightly delta-positive, given that it is a function of the performance of the stock since inception.

Orderly dispersion trading

As we pointed out before, dispersion is calculated by the difference between the realized index volatility and the market-cap weighted sum of the realized volatility of its constituents. The dispersion between time 0 and T is thus described by:

$$Dispersion_t = \sum_{i=1}^N \alpha_{i,t} \sigma_{i,t}^2 - \sigma_{I,t}^2$$

Where $\alpha_{i,t}$ is the weight of stock i in the index. By definition, it changes over time as stock prices change and is equal to:

$$\alpha_{i,t} = n_i S_{i,t} / I_t$$

If an investor trades dispersion using variance swaps weighted by the initial weights of the stocks in the index, he faces the risk of a possible change in weights over time until maturity of the variance swap. Gamma swaps, however, offer a more efficient way to trade dispersion. If one sells a gamma swap on the index and buys $n_i S_{i,0} / I_0$ – or $\alpha_{i,0}$ – gamma swaps on each stock i , the payoff at maturity should be equal to:

$$P \& L_T = \sum_{i=1}^N \frac{n_i S_{i,0}}{I_0} \frac{1}{T} \int_0^T \sigma_{i,t}^2 S_{i,t} / S_{i,0} dt - \frac{1}{T} \int_0^T \sigma_{I,t}^2 I_t / I_0 dt$$

Rearranging the terms, it gives:

$$\begin{aligned} P \& L_T &= \frac{1}{T} \int_0^T \left(I_t / I_0 \sum_{i=1}^N \alpha_{i,t} \sigma_{i,t}^2 \right) dt - \frac{1}{T} \int_0^T \sigma_{I,t}^2 I_t / I_0 dt \\ &= \frac{1}{T} \int_0^T \frac{I_t}{I_0} Dispersion_t dt \end{aligned}$$

The payoff is thus equal to the average dispersion over the period [0,T] weighted by index performance.

Forward-start variance swaps

Forward-start variance swaps are variance swaps whose variance is calculated between two future dates T and T'. The pay-off at maturity T' is therefore equal to:

$$\sigma^2 = \frac{252}{T' - T} \sum_{i=T+1}^{T'} (\ln(S_i / S_{i-1}))^2$$

Gamma swaps are a means to optimally trade dispersion.

Forward-start variance swaps enable positions to be taken in variance between two future dates.

Or in continuous-time:

$$\hat{\sigma}_{T,T'}^2 = \frac{1}{T-T'} \int_T^{T'} \sigma_t^2 dt$$

Carr and Madan again show that forward-start variance swaps may also be priced as the difference between a variance swap maturing at T' and a variance swap maturing at T . The value of a forward-start variance swap maturing at T' and starting on T is thus equal to:

$$V_f[\sigma_{T,T'}^2] = \frac{2e^{rT'}}{T'} \left[\int_0^{F_0} \frac{1}{K^2} P_0(K, T') dK + \int_{F_0}^{\infty} \frac{1}{K^2} C_0(K, T') dK \right] - \frac{2e^{rT}}{T} \left[\int_0^{F_0} \frac{1}{K^2} P_0(K, T) dK + \int_{F_0}^{\infty} \frac{1}{K^2} C_0(K, T) dK \right]$$

As a result, forward-start variance swaps enable investors to take positions on future volatility without having to enter two different variance swaps.

Corridor variance swaps

Corridor variance swaps are a variant of variance swaps that only take into account daily stock variations when the stock is in a specific range.

The payoff is equal to:

$$\text{Payoff}_{\text{Corridor}} = K^2 - K_{\text{Corr}}^2$$

where K^2 is the initial strike price and K_{Corr}^2 is described by:

$$K_{\text{Corr}}^2 = \frac{252}{T} \sum_{i=1}^T (\ln(S_i/S_{i-1}))^2 1_{S_{i-1} \in [K-\Delta; K+\Delta]}$$

Hence, squared returns are counted in if the stock price lies within a pre-specified range $[K-\Delta; K+\Delta]$.

Corridor variance swaps therefore enable bets to be taken on the pattern of the stock. If the stock move sideways and stays within the defined range, K_{Corr}^2 will be high. If the stock moves sharply upward or downward and leaves the range quickly, K_{Corr}^2 will be low.

Up and down corridor variance swaps

Up and down corridor variance swaps allow for betting on volatility on the upside or downside.

Financial engineering never stops producing new products and up corridor variance swaps are an example of such innovation. Up corridor variance swaps are a variant of corridor variance swaps and have the following payoff:

$$\text{Payoff}_{\text{Up}} = K^2 - K_{\text{UpCorr}}^2$$

Where K^2 is the initial strike price and K_{UpCorr}^2 is described by:

$$K_{\text{UpCorr}}^2 = \frac{252}{T} \sum_{i=1}^T (\ln(S_i/S_{i-1}))^2 1_{S_{i-1} > B}$$

Hence, squared returns are counted in if and only if the stock price lies above a predefined level denoted B . Indeed, one can also define a down corridor variance swap whose payoff would be defined by:


$$K^2_{downcorr} = \frac{252}{T} \sum_{i=1}^T (\ln(S_i/S_{i-1}))^2 1_{S_{i-1} < B}$$

Aggregating a down corridor variance swap and an up corridor variance swap yields the classic variance swap. This particular payoff has several advantages:

- It enables investors to bet on volatility or to hedge a position up to a certain level that may be defined by the overall risk structure of a derivatives book
- It is cheaper than a variance swap as only returns associated with a stock higher than B will be taken into account

Another reason why it is cheaper is that no purchases of expensive OTM puts are required to replicate it. The following graphs provide a comparison of the payoffs of 1-year variance swaps and 1-year ATM down corridor variance swaps on the S&P 500 since 1976. The ATM down corridor variance swap has a threshold level B equal to the initial level of the index, while only returns associated with an index level below the threshold are counted in.

Comparison of variance swaps and down corridor variance swaps on the S&P 500 index



Source: BNP Paribas

The ATM down corridor variance swap systematically yields a lower payoff and should thus be cheaper than the variance swap. Furthermore, the graph shows that the down corridor variance swap's payoff is negatively correlated with the market trend. With squared returns not counted in when the market rises, the down corridor variance swap offers a lower strike price. This is confirmed by the statistics reported in the following table. As a result, up and down corridor variance swaps enable a combined view to be taken on volatility, correlation and market direction.

Comparison of 1-year variance swaps with 1-year ATM up and down corridor variance swaps (in vol terms)

	Variance Swap	up corridor variance swap	down corridor variance swap	variance - up corridor	variance - down corridor
Average	14.91	10.69	8.10	4.22	6.81
Minimum	7.56	0.00	0.00	0.00	0.00
Maximum	36.84	32.92	36.82	36.82	27.14
correlation with 1-year S&P 500 returns	-31%	50%	-69%	-73%	64%

Source: BNP Paribas, CSFB/Tremont

Up and down conditional variance swaps

Up conditional variance swaps are a variant of up corridor variance swaps, whose payoff is described by:

$$Payoff_{Upcond} = (K^2 - K_{Upcond}^2) \sum_{i=1}^T 1_{S_{i-1} > B}$$

where K^2 is the initial strike price and K_{Upcond}^2 is described by:

$$K_{Upcond}^2 = \frac{252}{\sum_{i=1}^T 1_{S_{i-1} > B}} \sum_{i=1}^T (\ln(S_i / S_{i-1}))^2 1_{S_{i-1} > B}$$

Similarly, a down conditional variance swap can be structured as:

$$Payoff_{Downcond} = (K^2 - K_{Downcond}^2) \sum_{i=1}^T 1_{S_{i-1} < B}$$

where K^2 is the initial strike price and $K_{Downcond}^2$ is described by:

$$K_{Downcond}^2 = \frac{252}{\sum_{i=1}^T 1_{S_{i-1} < B}} \sum_{i=1}^T (\ln(S_i / S_{i-1}))^2 1_{S_{i-1} < B}$$

The conditional variance swap payoff is such that if the stock price never trades above the threshold B , it will be null whereas the up corridor variance swap payoff would be equal to K^2 . It also enables bets to be taken on very specific volatility behavior. Take the case of an index whose initial value is 100. If the index stands above 100 for a few days with a high volatility, say 40% annualized, and then drops below 100, an ATM up conditional variance swap will yield a much higher payoff than an ATM up corridor variance swap, given that the payoff's floating leg is divided by the number of days the index stays above the threshold (above 100 in this case). As a result, timing is less of an issue for an investor who buys a conditional variance swap rather than a corridor variance swap.

Riding the smile

One of the primary features of implied volatility is the presence of a "smile", i.e. the fact that implied volatilities depend on the option's strike. The goal of this section is not to explain why the smile exists, but how to trade it.

Optimally, one would need to buy or sell corridor variance swaps whose range $[k-\Delta, k+\Delta]$ would be such that Δ tends to zero. In this case, the conditional

New derivative products provide ways to take positions on implied volatility smiles.


variance swap would pay when the index is around k . Although theoretical models exist for pricing such products, they are not traded as yet¹².

Nevertheless, investors can use other methods based on existing volatility products to take views on the smile. The first one involves buying a variance swap and selling a gamma swap. As shown previously, variance swaps and gamma swaps can be replicated by a portfolio of options comprising a continuum of strike prices, inversely weighted by the square of the strike price and the strike price, respectively. This means that the spread between the variance swap and gamma swap strike prices should be an increasing function of the skew.

The following graph displays the smile for three different (90%-110%) skews, but which display the same average implied volatility of 22.5%. The table plots the prices of 1-year variance swaps and gamma swaps for the three different skews. One can see that the spread between the two increases when the skew rises. Indeed, variance swaps and gamma swaps yield different vega exposures. By being long a variance swap and short a gamma swap, an investor does not get a vega-neutral position. As a result, trading the smile by this method further requires to delta-hedge the position in order to cancel the vega sensitivity of the strategy. As delta-hedging necessitates an option model, investors further face model risk.

Note that gamma swaps generally yield lower payoffs than variance swaps, but that trading variance swaps against gamma swaps is a simple technique for trading the smile or the skew.

¹² See again Peter Carr and Dilip Madan, 1998, 'Towards a Theory of Volatility Trading', in *Volatility: New Estimation Techniques for Pricing Derivatives*, pp. 417-427, for the pricing of conditional variance swaps along a specific strike.

Indicative strike prices of 1-year gamma and variance swaps as a function of skew (90% - 110%)


	variance swap	gamma swap	spread
flat skew	21.30%	21.44%	-0.14%
4% skew	22.64%	21.81%	0.83%
6% skew	24.05%	22.63%	1.42%

Source – BNP Paribas

Another type of smile trade involves buying an ATM up-conditional variance-swap and selling an ATM down-conditional variance-swap. If an investor considers the smile to be relatively steep, i.e. if the skew is high, the up-conditional variance-swap should be cheap while the down conditional variance-swap should be expensive.

Correlation trading

Correlation measures how closely stocks move together or in opposite directions. The correlation between two stocks is denoted by $\rho_{i,j}$ and calculated as:

$$\rho_{i,j} = \frac{\text{cov}(R_i, R_j)}{\sigma_i \sigma_j} = \frac{\sum_{t=1}^T R_{i,t} R_{j,t}}{\sqrt{\sum_{t=1}^T R_{i,t}^2} \sqrt{\sum_{t=1}^T R_{j,t}^2}}$$

Correlation is an important factor as it strongly impacts index volatility. By definition, index volatility can be broken down as a function of the volatilities and pairwise correlations of its constituents:

$$\sigma_I^2 = \sum_{i=1}^n \alpha_i^2 \sigma_i^2 + \sum_{i=1}^{n-1} \sum_{j=i}^n \alpha_i \alpha_j \sigma_i \sigma_j \rho_{i,j}$$

The higher the correlation the higher the index volatility, since both the index weights α_i and the stock volatilities σ_i are positive. Average correlation provides a good measure of overall correlation. Average correlation is calculated by assuming the same pairwise correlation for all stocks:

$$\sigma_I^2 = \sum_{i=1}^n \alpha_i^2 \sigma_i^2 + \sum_{i=1}^{n-1} \sum_{j=i}^n \alpha_i \alpha_j \sigma_i \sigma_j \rho$$

Correlation is easily tradeable using correlation swaps.

$$\rho = \frac{\sigma_I^2 - \sum_{i=1}^n \alpha_i^2 \sigma_i^2}{\sum_{i=1}^{n-1} \sum_{j=i}^n \alpha_i \alpha_j \sigma_i \sigma_j}$$


The same breakdown holds for both implied and realized volatilities, thus meaning that option prices can be used to extract an average implied correlation for a given maturity.

There are several interesting features with correlation:

■ **Boundary and mean-reversion:**

Although by definition, correlation ranges from -1 to 1, it is usually positive as stocks tend to co-vary positively together. In view of this, correlation has a tendency to revert back towards its long-term mean.

6-week implied and realized correlation for the €Stoxx 50



Source – BNP Paribas

■ **Stock market and correlation:**

The fact that stocks generally all tend to head down when the bears are out means that correlation usually rises when markets decline.

Trading correlation can take different forms. One can approximate future correlation using variance or gamma swaps but the replication will never be perfect¹³. However, the simplest way to trade correlation is to take a position in a correlation swap.

A correlation swap has a payoff similar to the variance swap and one which is equal to:

$$Payoff_{correl} = (\rho_{realized} - \rho_{Strike}) * N$$

¹³ see N. Blanc, Index Variance Arbitrage: Arbitrating Component Correlation, BNP Paribas technical studies, 2004, for a discussion regarding several methods for optimizing correlation trades.

$\rho_{realized}$ is the average of all correlations $\rho_{i,j}$. $\rho_{i,j}$ is the correlation of the log daily returns of the stocks prices between stock i and j :

$$\rho_{Realised} = \frac{2}{n \times (n - 1)} \sum_{i=1}^{n-1} \sum_{j=i}^n \rho_{i,j}$$

where n is the number of stocks in the index or basket.

Conclusion

The combination of greater liquidity in the options market and new quantitative tools has now made trading volatility almost as easy as trading stocks or bonds. Variance swaps are now traded every business day on both indices and single stocks.

Financial engineering has created a wide range of new volatility products in addition to variance swaps. Derivative products are now available for investors to bet asymmetrically on volatility, on correlation, on the smile etc.

These financial instruments are not just arbitrage tools, but also efficient hedging tools. As we have shown in this report, they can be used to hedge structured product books or to hedge volatility up to a certain level.

In conclusion, every investor will probably find a volatility-based product that answers his needs, while new volatility instruments are bound to emerge in the future.

Appendix

Hedging options when volatility is unknown

In this section, we show how to break down the P&L of hedged options. We assume frictionless markets, no arbitrage and a constant risk-free interest rate. We also assume that the dynamics of the stock price are defined by the following geometric Brownian motion:

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_t dW_t$$

Where μ_t is the stock return, σ_t its unknown instantaneous volatility and W_t a Brownian motion. Let us denote the initial implied volatility by σ_i and assume that the trader chooses to hedge his option position by applying the Black-Scholes formula and using the hedge volatility σ_h . Under these assumptions, Carr shows that the P&L of an option V should be equal to:

$$P \& L = [V(S_0, 0; \sigma_i) - V(S_0, 0; \sigma_h)]e^{rT} + \int_0^T e^{r(T-t)} (\sigma_h^2 - \sigma_t^2) \frac{S_t^2}{2} \frac{\partial^2}{\partial S^2} V(S_t, t; \sigma_h) dt$$

Let us now follow Blanc's suggestion¹⁴ to linearize the first two terms in σ^2 around σ_h . The P&L can be rewritten as:

$$P \& L = Te^{rT} (\sigma_i^2 - \sigma_h^2) \frac{S_0^2}{2} \frac{\partial^2}{\partial S^2} V(S_0, 0; \sigma_h) + \int_0^T e^{r(T-t)} (\sigma_h^2 - \sigma_t^2) \frac{S_t^2}{2} \frac{\partial^2}{\partial S^2} V(S_t, t; \sigma_h) dt$$

Let us denote the volatility over the period 0 to T by $\hat{\sigma}$. $\hat{\sigma}$ is defined by:

$$\hat{\sigma}^2 = \frac{1}{T} \int_0^T \sigma_t^2 dt$$

The P&L may therefore be expressed as:

$$\begin{aligned} P \& L = & (\sigma_i^2 - \hat{\sigma}^2) T e^{rT} \frac{S_0^2}{2} \frac{\partial^2}{\partial S^2} V(S_0, T; \sigma_h) \\ & + (\hat{\sigma}^2 - \sigma_h^2) T \left[e^{rT} \frac{S_0^2}{2} \frac{\partial^2}{\partial S^2} V(S_0, T; \sigma_h) - \frac{1}{T} \int_0^T e^{r(T-t)} \frac{S_t^2}{2} \frac{\partial^2}{\partial S^2} V(S_t, T-t; \sigma_h) dt \right] \\ & + \int_0^T (\hat{\sigma}^2 - \sigma_t^2) e^{r(T-t)} \frac{S_t^2}{2} \frac{\partial^2}{\partial S^2} V(S_t, T-t; \sigma_h) dt \end{aligned}$$

Defining g_t as:

$$g_t = e^{r(T-t)} \frac{S_t^2}{2} \frac{\partial^2}{\partial S^2} V(S_t, T-t; \sigma_h) = \frac{e^{r(T-t)}}{2\sigma_h} \frac{\partial}{\partial \sigma} V(S_t, T-t; \sigma_h)$$

The P&L is thus equal to:

$$P \& L = (\sigma_i^2 - \hat{\sigma}^2) T g_0 + (\hat{\sigma}^2 - \sigma_h^2) T \left[g_0 - \frac{1}{T} \int_0^T g_t dt \right] + \int_0^T (\hat{\sigma}^2 - \sigma_t^2) g_t dt$$

¹⁴ See N. Blanc, Index Variance Arbitrage: Arbitrating Component Correlation, BNP Paribas technical studies, 2004.
https://eqd.bnpparibas.com/research/getdocument.asp?doc_id=126089&lang_id=1

Valuing gamma swaps

Carr and Madan's methodology for pricing variance swaps can also be applied to the valuation of gamma swaps. The continuous-time pay-off of a gamma swap is by definition equal to:

$$\Gamma_{0,T} = \int_0^T \sigma_t^2 S_t / S_0 dt$$

If $f(F_t, t)$ is a function of both F_t and t , then Ito's Lemma states that:

$$f(F_T) = f(F_0) + \int_0^T f'_F(F_t) dF_t + \int_0^T f'_t(F_t) dt + \int_0^T \frac{F_t}{2} f''(F_t) \sigma_t^2 dt$$

Let us consider the following function:

$$f(F_t, t) = e^{rt} (F_t \ln(F_t / F_0) - F_t + F_0)$$

Applying Ito's Lemma to $f(F_t, t)$ yields:

$$\frac{1}{2} \int_0^T S_t \sigma_t^2 dt = e^{rT} [F_T \ln(F_T / F_0) - F_T + F_0] - r \int_0^T e^{rt} ((1 + F_t) \ln(F_t / F_0) - F_t + F_0) dF_t$$

Note that Carr and Madan show that if a market exists for futures options of all strikes, any payoff $f(F_T)$ of the futures price F_T can be broken down as:

$$f(F_T) = f(\kappa) + f'(\kappa) [(F_T - \kappa)^+ - (\kappa - F_T)^+] + \int_0^\kappa f''(\kappa) (\kappa - F_T)^+ d\kappa + \int_\kappa^\infty f''(\kappa) (F_T - \kappa)^+ d\kappa \quad (\text{A})$$

In the absence of arbitrage, the above breakdown must prevail among initial values. Therefore, the initial value of the payoff is equal to:

$$V_f^0 = f(\kappa) B_0 + f'(\kappa) [C_0(\kappa) - P_0(\kappa)] + \int_0^\kappa f''(\kappa) P_0(\kappa) d\kappa + \int_\kappa^\infty f''(\kappa) C_0(\kappa) d\kappa$$

Applying equation (A) to the function $f(F_t, t) = F_t \ln(F_t / F_0) - F_t + F_0$ yields:

$$F_T \ln(F_T / F_0) - F_T + F_0 = \int_0^{F_0} \frac{1}{K} (K - F_T)^+ dK + \int_{F_0}^\infty \frac{1}{K} (F_T - K)^+ dK$$

Therefore, in order to receive $\int_0^T S_t \sigma_t^2 dt$ at time T , a trader should buy a continuum of puts with strike prices ranging from 0 to F_0 and calls with strikes ranging from F_0 to infinity, with everything weighted by the price of a discount bond e^{rT} . The initial cost is equal to:

$$e^{rT} \left[\int_0^{F_0} \frac{2}{K} P_0(K) dK + \int_{F_0}^\infty \frac{2}{K} C_0(K) dK \right] \quad (\text{B})$$

The trader also needs to roll a futures position, holding at t :

$$-2r((1 + F_t) \ln(F_t / F_0) - F_t + F_0) \quad (\text{C})$$

The net payoff of (B) and (C) at maturity T is:

$$\begin{aligned}
 & e^{rT} \int_0^{F_0} \frac{2}{K} (K - F_T)^+ dK + \int_{F_0}^{\infty} \frac{2}{K} (F_T - K)^+ dK - 2 \int_0^T re^{rt} \left((1 + F_t) \ln \left(\frac{F_t}{F_0} \right) - F_t + F_0 \right) dF_t \\
 &= 2e^{rT} \left(F_T \ln \left(\frac{F_T}{F_0} \right) - F_T + F_0 \right) - 2 \int_0^T re^{rt} \left((1 + F_t) \ln \left(\frac{F_t}{F_0} \right) - F_t + F_0 \right) dF_t \\
 &= \int_0^T S_t \sigma_t^2 dt
 \end{aligned}$$

Since the initial cost of executing this strategy is given by (B), the fair forward value of the gamma swap at time 0 should be equal to:

$$V_f[\Gamma_{0,T}] = \frac{2e^{2rT}}{TS_0} \left[\int_0^{F_0} \frac{1}{K} P_0(K) dK + \int_{F_0}^{\infty} \frac{1}{K} C_0(K) dK \right]$$

The Derman *et al.* replication of variance swaps and its extension to gamma swaps

Derman *et al.* suggest a simple discrete approximation of variance swaps using a fixed set of options. We previously showed that variance could be replicated by a log contract of the form


$$f(F_T) = \frac{2}{T} \left[\ln(F_0/F_T) + \frac{F_T}{F_0} - 1 \right]$$

plus a dynamic position on a futures contract. We also showed that the initial value of this log contract should be equal to:

$$\frac{2}{T} \int_0^{F_0} \frac{1}{K^2} P_0(K) dK + \int_{F_0}^{\infty} \frac{1}{K^2} C_0(K) dK$$

The issue for valuing the above continuum of options is that in practice, only a limited number of options are available. Therefore, Derman *et al.* propose a piecewise linear approximation of $f()$ as described by the following graph.

Log payoff and its discrete approximation in the case of variance swaps



Source – BNP Paribas

The segment between K_0 and K_{1C} resembles the payoff of a call with strike K_0 and the number of options one needs to buy is equal to the slope of the segment:

$$w_c(K_0) = \frac{f(K_{1C}) - f(K_0)}{K_{1C} - K_0}$$

The segment $[K_{1C}, K_{2C}]$ resembles a combination of calls with strike prices K_0 and K_{1C} , bearing in mind that we already own w_c calls with strike price K_0 . The number of options with strike price K_1 should thus be equal to:

$$w_c(K_1) = \frac{f(K_{2C}) - f(K_{1C})}{K_{2C} - K_{1C}} - w_c(K_0)$$

As a result, this method may be used to calculate all the weights for the calls position as well as the number of puts needed to be bought in order to replicate the left-hand side of the log payoff.

Indeed, this method can be also applied to the valuation of gamma swaps, the only difference being the function $f()$ we use. When valuing gamma swaps, the function f is equal to:

$$f(F_T) = \frac{2e^{rT}}{T} (F_T \ln(F_T/F_0) - F_T + F_0)$$

Example of capped variance-swap termsheet

VARIANCE SWAP BUYER	BNP PARIBAS
VARIANCE SWAP SELLER	CLIENT
UNDERLYING STOCK	ABC Corporation (Bloomberg Ticker: ABC <Equity>)
CURRENCY	USD
APPROXIMATE VEGA NOTIONAL	20,000
VARIANCE UNITS	294.1176, determined on the basis of the Approximate Vega Notional divided by 2 times the Volatility Strike.
TRADE DATE	August 11, 2005
EFFECTIVE DATE	August 11, 2005
FINAL VALUATION DATE	September 16, 2005
VALUATION DATES	Each Exchange Business Day from and including the Effective Date to and including the Final Valuation Date, regardless of the occurrence of a Market Disruption Event.
VOLATILITY STRIKE	34.00
VARIANCE STRIKE	$(34.00)^2 = 1,256.00$
INITIAL STOCK LEVEL (P ₁)	Closing Level of the Stock on Trade Date
EQUITY AMOUNT	<p>Variance Swap Buyer will pay:</p> $\text{Variance Units} * \text{Max} [0, \text{Variance Strike} - \text{Min}((2.5 * \text{Volatility strike})^2; \text{Realized Volatility}^2)]$ <p>Variance Swap Seller will pay:</p> $\text{Variance Units} * \text{Max}[0, \text{Min}((2.5 * \text{Volatility strike})^2; \text{Realized Volatility}^2) - \text{Variance Strike}]$

REALIZED VOLATILITY

$$100X \sqrt{\frac{\sum_{i=1}^{n-1} \left[\ln \left(\frac{(P_{i+1}) + (D_{i+1})}{P_i} \right) \right]^2}{n-1}} X \sqrt{\text{Business Days Per Year}}$$

where:

n = number of expected Valuation Dates, known as of trade date

P_i = the Closing Level on the Exchange of the Stock on the i^{th} Exchange Business Day from and including the Effective Date to and including the Final Valuation Date ($i=n$)

P_n = shall be equal to the Closing Level on the Exchange of the Stock on the Final Valuation Date

Business Days Per Year = [252]

If there is any Dividend $D(i)$, ex-date at date i , the value of the observation calculated for day I should be $P(i) + D(i)$ where $P(i)$ is the closing price and the observed return is $(P(i) + D(i)) / P(i-1)$.

Should a Market Disruption Event occur on any Exchange Business Day, then the Closing Price for that day will equal the Closing Price on the first preceding Exchange Business Day on which there was no Market Disruption Event.

CASH SETTLEMENT PAYMENT DATE	Three Currency Business Days following the Final Valuation Date
CALCULATION AGENT	BNP Paribas and the Counterparty
GOVERNING LAW	New York
INDEPENDENT AMOUNT	Not Applicable
DOCUMENTATION	If the parties are not parties to an ISDA Master Agreement covering this Transaction, the Counterparty and BNP Paribas will enter into a master agreement together with a Credit Support Annex and this Transaction will supplement, form a part of, and be subject to such modifications as the parties will agree in good faith. The Confirmation will incorporate the 2002 Equity Derivatives Definitions and 2000 ISDA Definitions.

References and further reading

Blanc, Nicolas, 2004, Index Variance Arbitrage: Arbitraging Component Correlation, BNP Paribas technical studies

Carr, Peter, FAQs in Option Pricing Theory, Journal of Derivatives, forthcoming

Carr, Peter and Dilip Madan, 1998, Towards a Theory of Volatility Trading, in Volatility: New Estimation Techniques for Pricing Derivatives, p. 417-427

Carr, Peter and Liuren Wu, 2005, Variance Risk Premia, Courant Institute working paper

Chance, Don M., 1995, A Chronology of Derivatives, Derivative Quarterly, p. 53-60

Derman, Emanuel, Michael Kamal, Joseph Zou and Kresimir Demeterfi, 1999, A Guide to Volatility and Variance Swaps, The Journal of Derivatives, Summer p.1-32

Driessen, Joost, Pascal Maenhout and Grigory Vilkov, 2005, Option-Implied Correlations and the Price of Correlation Risk, INSEAD Working Paper

In accordance with US regulations, the analyst(s) named in this report certifies that 1) all of the views expressed herein accurately reflect the personal views of the analyst(s) with regard to any and all of the content, securities and issuers and 2) no part of the compensation of the analyst(s) was, is, or will be, directly or indirectly, related to the specific recommendation or views expressed by the analyst(s) in this report.

It is a BNP Paribas policy that analysts are prohibited from trading in any of the stocks in the sector that they cover.

THIS PUBLICATION IS NOT INTENDED FOR PRIVATE CUSTOMERS AS DEFINED IN THE FSA RULES AND SHOULD NOT BE PASSED ON TO ANY SUCH PERSON.

The material in this report was produced by BNP Paribas SA, a limited company, whose head office is in Paris. BNP Paribas SA is registered as a bank with the "Comité des Etablissements de Crédit et des Entreprises d'Investissement ("CECEI") and is regulated by the "Autorité des Marchés Financiers ("AMF") for the conduct of its designated investment business in France.

The information and opinions contained in this report have been obtained from public sources believed to be reliable, but no representation or warranty, express or implied, is made that such information is accurate or complete and it should not be relied upon as such. 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, and are subject to change without notice. This report is not, and should not be construed as, an offer document or an offer or solicitation to buy or sell any investments. Any reference to past performance should not be taken as an indication of future performance. No BNP Paribas Group Company accepts any liability whatsoever for any direct or consequential loss arising from any use of material contained in this report. This report is confidential and is submitted to selected recipients only. It may not be reproduced (in whole or in part) to any other person. A BNP Paribas Group Company and/or persons connected with it may effect or have effected a transaction for their own account in the investments referred to in the material contained in this report or any related investment before the material is published to any BNP Paribas Group Company's customers. On the date of this report a BNP Paribas Group Company, persons connected with it and their respective directors and/or representatives and/or employees may have a long or short position in any of the investments mentioned in this report and may purchase and/or sell the investments at any time in the open market or otherwise, in each case either as principal or as agent. Additionally, a BNP Paribas Group Company within the previous twelve months may have acted as an investment banker or may have provided significant advice or investment services to the companies or in relation to the investment(s) mentioned in this report.

This report is prepared for professional investors and is not intended for Private Customers, and should not be passed on to any such persons. For the purpose of distribution in the United States this report is only intended for persons which can be defined as 'Major Institutional Investors' under U.S. regulations. Any U.S. person receiving this report and wishing to effect a transaction in any security discussed herein, must do so through a U.S. registered broker dealer. BNP Paribas Securities Corp. is a U.S. registered broker dealer.

By accepting this document you agree to be bound by the foregoing limitations.

© BNP Paribas (2004). All rights reserved.