JPMorgan

JUST WHAT YOU NEED TO KNOW ABOUT VARIANCE SWAPS

Sebastien Bossu

Eva Strasser

Regis Guichard

Equity Derivatives Investor Marketing

Quantitative Research & Development

JPMorgan – London

EQUITY DERIVATIVES

JPMorgan

Overview

In this note we introduce the properties of variance swaps, and give details on the hedging and valuation of these instruments.

- Section 1 gives quick facts about variance swaps and their applications.

- Section 2 is written for traders and market professionals who have some degree of familiarity with the theory of vanilla option pricing and hedging, and explains in 'intuitive' mathematical terms how variance swaps are hedged and priced.

- Section 3 is written for quantitative traders, researchers and financial engineers, and gives theoretical insights into hedging strategies, impact of dividends and jumps.

- Appendix A is a review of the concepts of historical and implied volatility.

- Appendices B and C cover technical results used in the note.

We thank Cyril Levy-Marchal, Jeremy Weiller, Manos Venardos, Peter Allen, Simone Russo for their help or comments in the preparation of this note.

JPMorgan

Table of Contents

JPMorgan

1. Variance Swaps

1.1. Payoff

A variance swap is an instrument which allows investors to trade future realized (or historical) volatility against current implied volatility. As explained later in this document, only variance —the squared volatility— can be replicated with a static hedge. [See Sections 2.2 and 3.2 for more details.]

Sample terms are given in Exhibit 1.1.1 below.

Exhibit 1.1.1 — Variance Swap on S&P 500 : sample terms and conditions

<div style="border:1px solid">

VARIANCE SWAP ON S&P500

SPX INDICATIVE TERMS AND CONDITIONS

Instrument:	Swap
Trade Date:	TBD
Observation Start Date:	TBD
Observation End Date:	TBD
Variance Buyer:	TBD (e.g. JPMorganChase)
Variance Seller:	TBD (e.g. Investor)
Denominated Currency:	USD ("USD")
Vega Amount:	100,000
Variance Amount:	3,125 (determined as Vega Amount/(Strike*2))
Underlying:	S&P500 (Bloomberg Ticker: SPX Index)
Strike Price:	**16**
Currency:	USD

Equity Amount: **T+3** after the Observation End Date, the Equity Amount will be calculated and paid in accordance with the following formula:

*Final Equity payment = Variance Amount * (Final Realized Volatility^2 – Strike Price^2)*

If the Equity Amount is positive the Variance Seller will pay the Variance Buyer the Equity Amount.
If the Equity Amount is negative the Variance Buyer will pay the Variance Seller an amount equal to the absolute value of the Equity Amount.

where

$$Final\ Realised\ Volatility = \sqrt{\frac{252 \times \sum_{t=1}^{t=N}\left(\ln\frac{P_t}{P_{t-1}}\right)^2}{Expected\_N}} \times 100$$

Expected_N = [number of days], being the number of days which, as of the Trade Date, are expected to be Scheduled Trading Days in the Observation Period
P_0 = The Official Closing of the underlying at the Observation Start Date
P_t = Either the Official Closing of the underlying in any observation date t or, at Observation End Date, the Official Settlement Price of the Exchange-Traded Contract

Calculation Agent:	JP Morgan Securities Ltd.
Documentation:	ISDA

</div>

Note:

- Returns are computed on a logarithmic basis: $\ln\left(\dfrac{P_t}{P_{t-1}}\right)$.

- The mean return, which normally appears in statistics textbooks, is dropped. This is because its impact on the price is negligible (the expected average daily return is $1/252^{nd}$ of the money-market rate), while its omission has the benefit of making the payoff perfectly additive (3-month variance + 9-month variance in 3 months = 1-year variance.)

- It is a market practice to define the variance notional in volatility terms:

$$Variance\ Notional = \frac{Vega\ Notional}{2 \times Strike}$$

With this adjustment, if the realized volatility is 1 'vega' (volatility point) above the strike at maturity, the payoff is approximately equal to the Vega Notional.

Convexity

The payoff of a variance swap is convex in volatility, as illustrated in Exhibit 1.1.2. This means that an investor who is long a variance swap (i.e. receiving realized variance and paying strike at maturity) will benefit from boosted gains and discounted losses. This bias has a cost reflected in a slightly higher strike than the 'fair' volatility[2], a phenomenon which is amplified when volatility skew is steep. Thus, the fair strike of a variance swap is often in line with the implied volatility of the 90% put.

Exhibit 1.1.2 — Variance swaps are convex in volatility

---
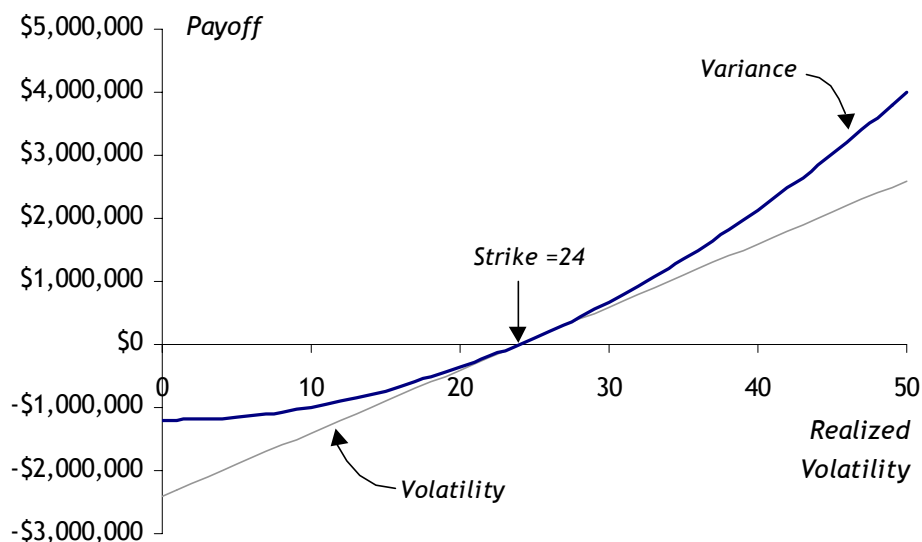
[2] Readers with a mathematical background will recall Jensen's inequality: $E(\sqrt{Variance}) \leq \sqrt{E(Variance)}$.

JPMorgan

Rules of thumb

Demeterfi—Derman—Kamal—Zou (1999) derived a rule of thumb for the fair strike of a variance swap when the skew is linear in strike:

$$K_{\text{var}} \approx \sigma_{ATMF} \sqrt{1 + 3T \times skew^2}$$

where σ_{ATMF} is the at-the-money-forward volatility, T is the maturity, and *skew* is the slope of the skew curve. For example, with σ_{ATMF} = 20%, T = 2 years, and a 90-100 skew of 2 vegas, we have $K_{\text{var}} \approx$ 22.3%, which is in line with the 90% put implied volatility normally observed in practice.

For log-linear skew, similar techniques give the rule of thumb:

$$K_{\text{var}} \approx \sqrt{\sigma_{ATMF}^2 + \beta \sigma_{ATMF}^3 T + \frac{\beta^2}{4}\left(12\sigma_{ATMF}^2 T + 5\sigma_{ATMF}^4 T^2\right)}$$

where σ_{ATMF} is the at-the-money-forward volatility, T is the maturity, and β is the slope of the log skew curve[3]. For example, with σ_{ATMF} = 20%, T = 2 years, and a 90-100 skew of 2 vegas, we have $\beta = -\dfrac{2\%}{\ln(0.9)} \approx 0.19$ and $K_{\text{var}} \approx$ 22.8%.

Note that these two rules of thumb produce good results only for non-steep skew.

1.2. Applications

Volatility Trading

Variance swaps are natural instruments for investors taking directional bets on volatility:

- Realized volatility: unlike the trading P&L of a delta-hedged option position, a long variance position will always benefit when realized volatility is higher than implied at inception, and conversely for a short position [see Section 2.1 on P&L path-dependency.]

- Implied volatility: similar to options, variance swaps are fully sensitive at inception to changes in implied volatility

Variance swaps are especially attractive to volatility sellers for the following two reasons:

- Implied volatility tends to be higher than final realized volatility: 'the derivative house has the statistical edge.'

- Convexity causes the strike to be around the 90% put implied volatility, which is slightly higher than 'fair' volatility.

Forward volatility trading

Because variance is additive, one can obtain a perfect exposure to forward implied volatility with a calendar spread. For example, a short 2-year vega exposure of €100,000 on the EuroStoxx 50 starting in 1 year can be hedged as follows [levels as of 21 April, 2005]:

---

[3] The skew curve is thus assumed to be of the form: $\sigma(K) = \sigma_{ATMF} - \beta \ln(K/F)$ where F is the forward price.

- Long 2-year variance struck at 19.50 on a Vega Notional of €200,000 (i.e. a Variance Notional of 5,128)
- Short 1-year variance struck at 18.50 on a Variance Notional of 5,128 / 2 = 2,564 (i.e. a Vega Notional of €94,868)

Implied forward volatility on this trade is approximately[4]:

$$\underbrace{19.50}_{2-\text{year vol}} \times \underbrace{2}_{\text{tenor}} - \underbrace{18.50}_{1-\text{year vol}} \times \underbrace{1}_{\text{tenor}} = 20.5.$$

Therefore, if the 1-year implied volatility is above 20.5 in one year's time, say at 21, the hedge will be approximately up ½ a vega, or €50,000, while the exposure will be down by the same amount.

However, keep in mind that the fair value of a variance swap is also sensitive to skew.

Forward volatility trades are interesting because the forward volatility term structure tends to flatten for longer forward-start dates, as illustrated in Exhibit 1.2.1 below. In this example, we can see that the 1-year forward volatilities exhibit a downard sloping term structure. Thus, an investor who believes that the term structure will revert to an upward sloping shape might want to sell the 12x1 and buy the 12x12 implied volatilities, or equivalently sell 13m and buy 24m, with appropriate notionals:

Buy 12x12	=	Buy 24m and Sell 12m
Sell 12x1	=	Sell 13m and Buy 12m
Buy spread	=	Buy 24m and Sell 13m

Exhibit 1.2.1 — Spot and forward volatility curves derived from fair variance swap strikes
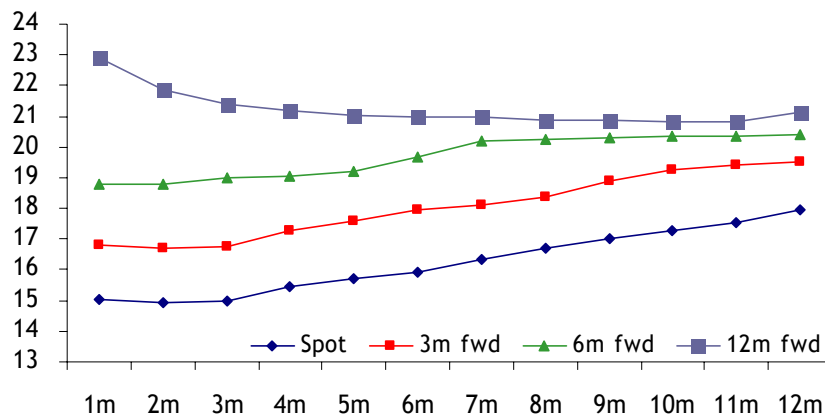
Source: JPMorgan.

Spreads on indices

Variance swaps can also be used to capture the volatility spread between two correlated indices, for instance by being long 3-month DAX variance and short 3-month EuroStoxx 50 variance. Exhibit 1.2.2 below shows that in the period 2000-2004 the historical spread was

---

[4] An accurate calculation would be: $\sqrt{2y\,vol^2 \times 2 - 1y\,vol^2 \times 1}$

almost always in favor of the DAX and sometimes as high as 12 vegas, while the implied spread[5] ranged between -4 and +4 vegas.

a)

b)
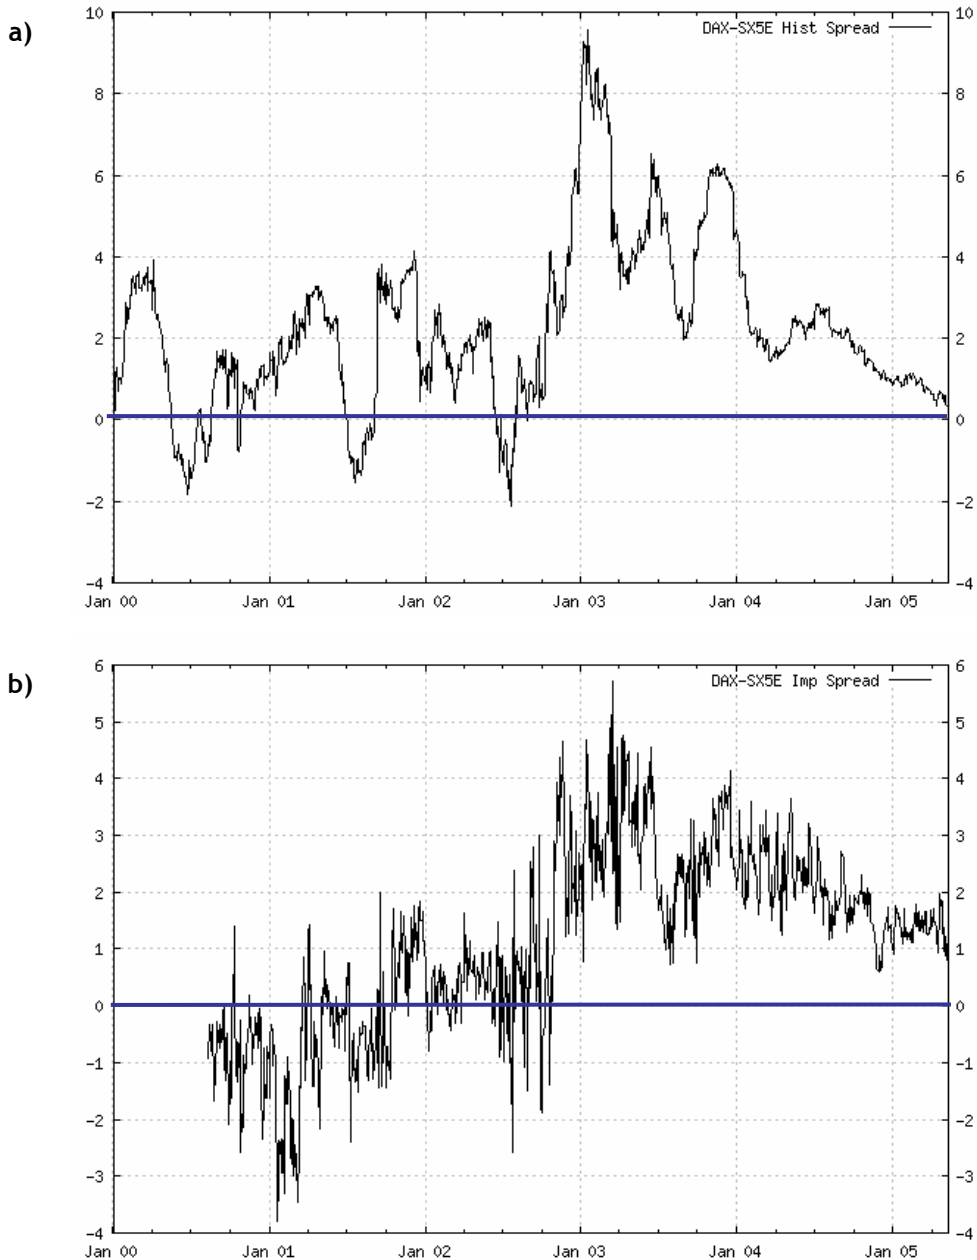
Source: JPMorgan—DataQuery.

Correlation trading: Dispersion trades

A popular trade in the variance swap universe is to sell correlation by taking a short position on index variance and a long position on the variance of the components. Exhibit 1.2.3 below shows the evolution of one-year implied and realized correlation.

---

[5] Measured as the difference between the 90% strike implied volatilities. Actual numbers may differ depending on skew, transaction costs and other market conditions.

JPMorgan

Exhibit 1.2.3 — Implied and realized correlation of EuroStoxx 50

Source: JPMorgan—DataQuery.

More formally the payoff of a variance dispersion trade is:

$$\sum_{i=1}^{n} w_i Notional_i \sigma_i^2 - Notional_{Index} \sigma_{Index}^2 - Residual\ Strike$$

where w's are the weights of the index components, σ's are realized volatilities, and notionals are expressed in variance terms. Typically, only the most liquid stocks are selected among the index components, and each variance notional is adjusted to match the same vega notional as the index in order to make the trade vega-neutral at inception.

1.3. Mark-to-market and Sensitivities

Mark-to-market

Because variance is additive in time dimension the mark-to-market of a variance swap can be decomposed at any point in time between realized and implied variance:

$$VarSwap_t = Notional \times PV_t(T) \times \left[\frac{t}{T} \times \left(Realized\ Vol(0,t) \right)^2 \right.$$

$$\left. + \frac{T-t}{T} \left(Implied\ Vol(t,T) \right)^2 - Strike^2 \right]$$

where Notional is in variance terms, $PV_t(T)$ is the present value at time t of \$1 received at maturity T, Realized Vol(0, t) is the realized volatility between inception and time t, Implied Vol(t, T) is the fair strike of a variance swap of maturity T issued at time t.

For example, consider a one-year variance swap issued 3 months ago on a vega notional of \$200,000, struck at 20. The 9-month zero-rate is 2%, realized volatility over the past 3 months

was 15, and a 9-month variance swap would strike today at 19. The mark-to-market of the one-year variance swap would be:

$$VarSwap_t = \frac{200,000}{2 \times 20} \times \frac{1}{(1+2\%)^{0.75}} \times \left[\frac{1}{4} \times 15^2 + \frac{3}{4} \times 19^2 - 20^2 \right]$$
$$= -\$359,619$$

Note that this is not too far from the 2 vega loss which one obtains by computing the weighted average of realized and implied volatility: 0.25 x 15 + 0.75 x 19 = 18, minus 20 strike.

Vega sensitivity

The sensitivity of a variance swap to implied volatility decreases linearly with time as a direct consequence of mark-to-market additivity:

$$Vega = \frac{\partial VarSwap_t}{\partial \sigma_{implied}} = Notional \times (2\sigma_{implied}) \times \frac{T-t}{T}$$

Note that Vega is equal to 1 at inception if the strike is fair and the notional is vega-adjusted:

$$Notional = \frac{Vega\ Notional}{2 \times Strike}$$

Skew sensitivity

As mentioned earlier the fair value of a variance swap is sensitive to skew: the steeper the skew the higher the fair value. Unfortunately there is no straightforward formula to measure skew sensitivity but we can have a rough idea using the rule of thumb for linear skew in Section 1.1:

$$K_{var}^2 \approx \sigma_{ATMF}^2 \left(1 + 3T \times skew^2 \right)$$

$$Skew\ Sensitivity \approx 6 \times Notional \times \sigma_{ATMF}^2 \times \frac{T-t}{T} \times skew$$

For example, consider a one-year variance swap on a vega notional of $200,000, struck at 15. At-the-money-forward volatility is 14, and the 90-100 skew is 2.5 vegas. According to the rule of thumb, the fair strike is approximately 14 x (1 + 3 x (2.5/10)^2) = 16.62. If the 90-100 skew steepens to 3 vegas the change in mark-to-market would be:

$$\Delta MTM \approx \underbrace{6 \times \frac{200,000}{2 \times 15} \times 14^2 \times \frac{2.5}{10}}_{Sensitivity} \times \underbrace{\left(\frac{3-2.5}{10} \right)}_{\Delta Skew} \approx \$100,000$$

Dividend sensitivity

Dividend payments affect the price of a stock, resulting in a higher variance. When dividends are paid at regular intervals, it can be shown that ex-dividend annualized variance should be

adjusted by approximately adding the square of the annualized dividend yield divided by the number of dividend payments per year[6]. The fair strike is thus:

$$K_{var} \approx \sqrt{(K_{var}^{ex-div})^2 + \frac{(Div\ Yield)^2}{Nb\ Divs\ Per\ Year}}$$

From this adjustment we can derive a rule of thumb for dividend sensitivity:

$$\mu = \frac{\partial VarSwap_t}{\partial Div\ Yield} \approx Notional \times \frac{Div\ Yield / Nb\ Divs\ Per\ Year}{K_{var}} \times \frac{T-t}{T}$$

For example, consider a one-year variance swap on a vega notional of $200,000 struck at 20. The fair strike ex-dividend is 20 and the annual dividend yield is 5%, paid semi-annually. The adjusted strike is thus $(20^2 + 5^2 / 2)^{0.5} = 20.31$. Were the dividend yield to increase to 5.5% the change in mark-to-market would be:

$$\Delta MTM \approx 200,000 \times \underbrace{\frac{5/2}{20.31}}_{skew\ sensitivity} \times \underbrace{(5.5-5)}_{\Delta Div\ Yield} \approx \$12,310$$

However, in the presence of skew, changes in dividend expectations will also impact the forward price of the underlying which in turns affects the fair value of varianc. This phenomenon will normally augment the overall dividend sensitivity of a variance swap.

---

[6] More specifically the adjustment is $\frac{1}{T}\sum_{j=1}^{M}\left(d_j \times \frac{T}{M}\right)^2 = D^2 \times \frac{T}{M}$ where d_1, d_2, ..., d_M are gross dividend yields and D is the annualized 'average' dividend yield. See Section 3.3 for more details.

JPMorgan

2. Valuation and Hedging in Practice

2.1. Vanilla Options: Delta-Hedging and P&L Path-Dependency

Delta-Hedging

Option markets are essentially driven by expectations of future volatility. This results from the way an option payoff can be dynamically replicated by only trading the underlying stock and cash, as described in 1973 by Black—Scholes and Merton.

More specifically, the sensitivity of an option price to changes in the stock price, or delta, can be entirely offset by continuously holding a reverse position in the underlying in quantity equal to the delta. For example, a long call position on the S&P 500 index with an initial delta of $5,000 per index point (worth $6,000,000 for an index level of 1,200) is delta-neutralized by selling 5,000 units of the S&P 500 (in practice 20 futures contracts: 6,000,000/(250 x 1,200)) Were the delta to increase to $5,250 per index point, the hedge should be adjusted by selling an additional 250 units (1 contract), and so forth. The iteration of this strategy until maturity is known as delta-hedging.

Once the delta is hedged, an option trader is mostly left with three sensitivities:

- Gamma: sensitivity of the option delta to changes in the underlying stock price ;

- Theta or time decay: sensitivity of the option price to the passage of time ;

- Vega: sensitivity of the option price to changes in the market's expectation of future volatility (i.e. implied volatility.)[7]

The daily P&L on a delta-neutral option position can be decomposed along these three factors:

$$\text{Daily P\&L = Gamma P\&L + Theta P\&L + Vega P\&L + Other} \qquad \textbf{(Eq. 1)}$$

Here 'Other' includes the P&L from financing the reverse delta position on the underlying, as well as the P&L due to changes in interest rates, dividend expectations, and high-order sensitivities (e.g. sensitivity of Vega to changes in stock price, etc.)

Equation 1 can be rewritten:

$$\text{Daily P\&L} = \frac{1}{2}\Gamma \times (\Delta S)^2 \quad + \quad \Theta \times (\Delta t) \quad + \quad \mathcal{V} \times (\Delta \sigma) \quad + ...$$

where ΔS is the change in the underlying stock price, Δt is the fraction of time elapsed (typically 1/365), and $\Delta \sigma$ is the change in implied volatility.

We now consider a world where implied volatility is constant, the riskless interest rate is zero, and other P&L factors are negligible. In this world resembling Black-Scholes, we have the reduced P&L equation:

$$\text{Daily P\&L} = \frac{1}{2}\Gamma \times (\Delta S)^2 \quad + \quad \Theta \times (\Delta t) \qquad \textbf{(Eq. 2)}$$

We proceed to interpret Equation 2 in terms of volatility, and we will see that in this world the daily P&L of a delta-hedged option position is essentially driven by realized and implied volatility.

---

[7] Note that in Black-Scholes volatility is assumed to remain constant through time. The concept of Vega is thus inconsistent with the theory, yet critical in practice.

We start with the well-known relationship between theta and gamma:

$$\Theta \approx -\frac{1}{2}\Gamma S^2 \sigma^2 \qquad \textbf{(Eq. 3)}$$

where S is the current spot price of the underlying stock and σ the current implied volatility of the option.

In our world with zero interest rate, this relationship is actually exact, not approximate. Appendix B presents two derivations of Equation 3, one based on intuition and one which is more rigorous.

Equation 3 is the core of Black-Scholes: it dictates how option prices diffuse in time in relation to convexity. Plugging Equation 3 into Equation 2 and factoring S^2, we obtain a characterization of the daily P&L in terms of squared return and squared implied volatility:

$$\text{Daily P\&L} = \frac{1}{2}\Gamma S^2 \times \left[\left(\frac{\Delta S}{S}\right)^2 - \sigma^2 \Delta t\right] \qquad \textbf{(Eq. 4)}$$

The first term in the bracket, $\frac{\Delta S}{S}$, is the percent change in the stock price — in other words, the one-day stock return. Squared, it can be interpreted as the *realized* one-day variance.

The second term in the bracket, $\sigma^2 \Delta t$, is the squared daily implied volatility, which one could name the daily implied variance.

Thus, Equation 4 tells us that the daily P&L of a delta-hedged option position is driven by the spread between realized and implied variance, and breaks even when the stock price movement exactly matches the market's expectation of volatility.

In the following paragraph we extend this analysis to the entire lifetime of the option.

P&L path-dependency

One can already see the connection between Equation 4 and variance swaps: if we sum all daily P&L's until the option's maturity, we obtain an expression for the *final* P&L:

$$\text{Final P\&L} = \frac{1}{2}\sum_{t=0}^{n}\gamma_t\left[r_t^2 - \sigma^2 \Delta t\right] \qquad \textbf{(Eq. 5)}$$

where the subscript t denotes time dependence, r_t the stock daily return at time t, and γ_t the option's gamma multiplied by the square of the stock price at time t, also known as **dollar gamma**.

Equation 5 is very close to the payoff of a variance swap: it is a weighted sum of squared realized returns minus a constant that has the role of the strike. The main difference is that in a variance swap weights are constant, whereas here the weights depend on the option gamma through time, a phenomenon which is known to option traders as the path-dependency of an option's trading P&L, illustrated in Exhibit 2.1.1.

It is interesting to note that even when the stock returns are assumed to follow a random walk with a volatility equal to σ, Equation 5 does not become nil. This is because each squared return remains *distributed* around $\sigma^2 \Delta t$ rather than equal to $\sigma^2 \Delta t$. However this particular

JPMorgan

path-dependency effect is mostly due to discrete hedging rather than a discrepancy between implied and realized volatility and will vanish in the case of continuous hedging[8].

In this example an option trader sold a 1-year call struck at 110% of the initial price on a notional of $10,000,000 for an implied volatility of 30%, and delta-heged his position daily. The realized volatility was 27.50%, yet his final trading P&L is down $150k. Furthermore, we can see (Figure a) that the P&L was up $250k until a month before expiry: how did the profits change into losses? One indication is that the stock price oscillated around the strike in the final months (Figure a), triggering the dollar gamma to soar (Figure b.) This would be good news if the volatility of the underlying remained below 30% but unfortunately this period coincided with a change in the volatility regime from 20% to 40% (Figure b.) Because the daily P&L of an option position is weighted by the gamma and the volatility spread between implied and realized was negative, the final P&L drowned, even though the realized volatility over the year was below 30%!

a)

Stock Price (Initial = 100)   'Hammered at the strike' !   Trading P&L ($)

b)

Stock Price (Initial = 100)   Volatility

---
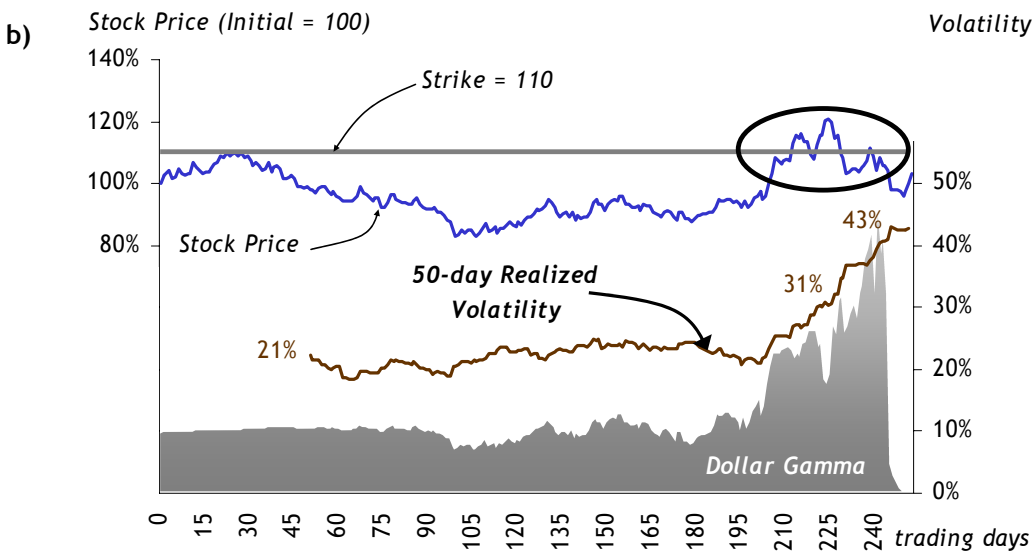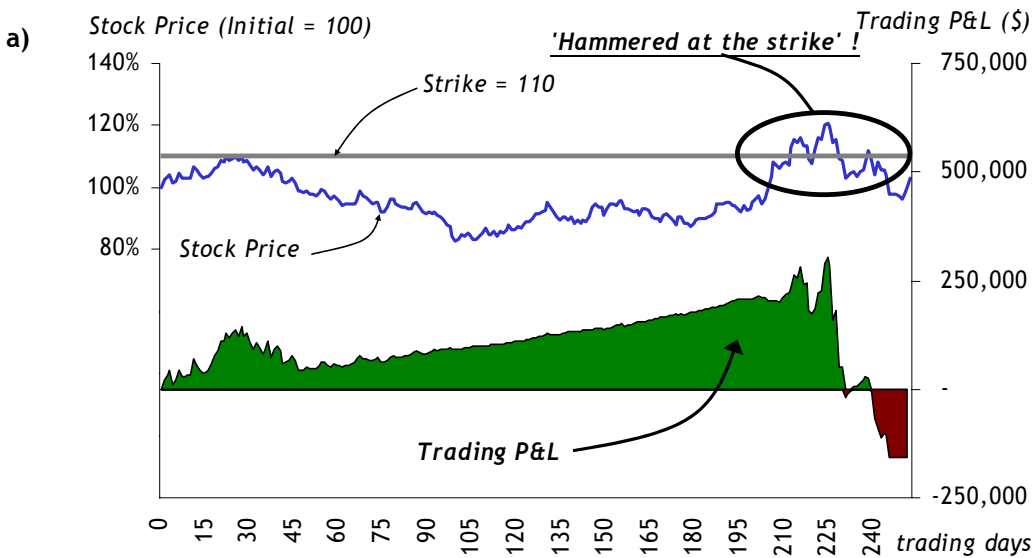
[8] See Wilmott (1998) for a theoretical approach of discrete hedging and Allen—Harris (2001) for a statistical analysis of this phenomenon. Wilmott notes that the daily Gamma P&L has a chi-square distribution, while Allen—Harris include a bell-shaped chart of the distribution of 1000 final P&Ls of a discretely delta-hedged option position. Neglecting the gamma dependence, the central-limit theorem indeed shows that the sum of n independent chi-square variables converges to a normal distribution.

JPMorgan

2.2. Static Replication of Variance Swaps

In the previous paragraph we saw that a vanilla option trader following a delta-hedging strategy is essentially replicating the payoff of a *weighted* variance swap where the daily squared returns are weighted by the option's dollar gamma[9]. We now proceed to derive a static hedge for standard ('non-gamma-weighted') variance swaps. The core idea here is to combine several options together in order to obtain a constant aggregate gamma.

Exhibit 2.2.1 shows the dollar gamma of options with various strikes in function of the underlying level. We can see that the contribution of low-strike options to the aggregate gamma is small compared to high-strike options. Therefore, a natural idea is to increase the weights of low-strike options and decrease the weights of high-strike options.

Dollar Gamma

Aggregate

K = 200
K = 175
K = 150
K = 125
K = 100
K = 75
K = 50
K = 25

0   25   50   75   100   125   150   175   200   225   250   275   300

Underlying Level (ATM = 100)
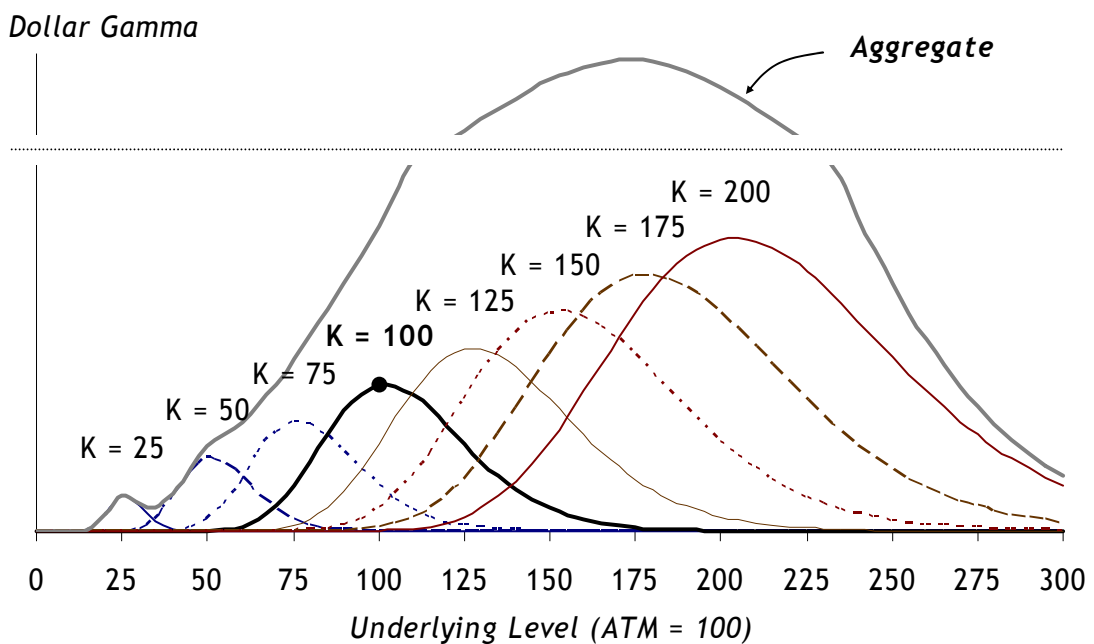
An initial, 'naïve' approach to this weighting problem is to determine individual weights $w(K)$ such that each option of strike K has a peak dollar gamma of, say, 100. Using the Black-Scholes closed-form formula for gamma, one would find that the weights should be inversely proportional to the strike (i.e. $w(K) = c / K$, where c is a constant.) [See Appendix C for details.]

Exhibit 2.2.2 shows the dollar gamma resulting from this weighting scheme. We can see that the aggregate gamma is still non-constant (whence the adjective 'naïve' to describe this approach), however we also notice the existence of a linear region when the underlying level is in the range 75—135.

---

[9] Recall that dollar gamma is defined as the second-order sensitivity of an option price to a percent change in the underlying. In this paragraph, we use the terms 'gamma' and 'dollar gamma' interchangeably.

This observation is crucial: if we can regionally obtain a linear aggregate gamma with a certain weighting scheme $w(K)$, then the modified weights $w'(K) = w(K)\ /\ K$ will produce a constant aggregate gamma. Since the naïve weights are inversely proportional to the strike K, the correct weights should be chosen to be inversely proportional to the *squared* strike, i.e.:

$$w(K) = \frac{c}{K^2}$$

where c is a constant.

Exhibit 2.2.3 shows the results of this approach for the individual and aggregate dollar gammas. As expected, we obtain a constant region when the underlying level stays in the range 75—135.

A perfect hedge with a constant aggregate gamma for all underlying levels would take infinitely many options struck along a continuum between 0 and infinity and weighted inversely proportional to the squared strike. This is etablished rigorously in Section 3.2. Note that this is a strong result, as the static hedge is both space (underlying level) and time independent.

Exhibit 2.2.3 — Dollar gamma of options weighted inversely proportional to the square of strike

Dollar Gamma

Underlying Level (ATM = 100)
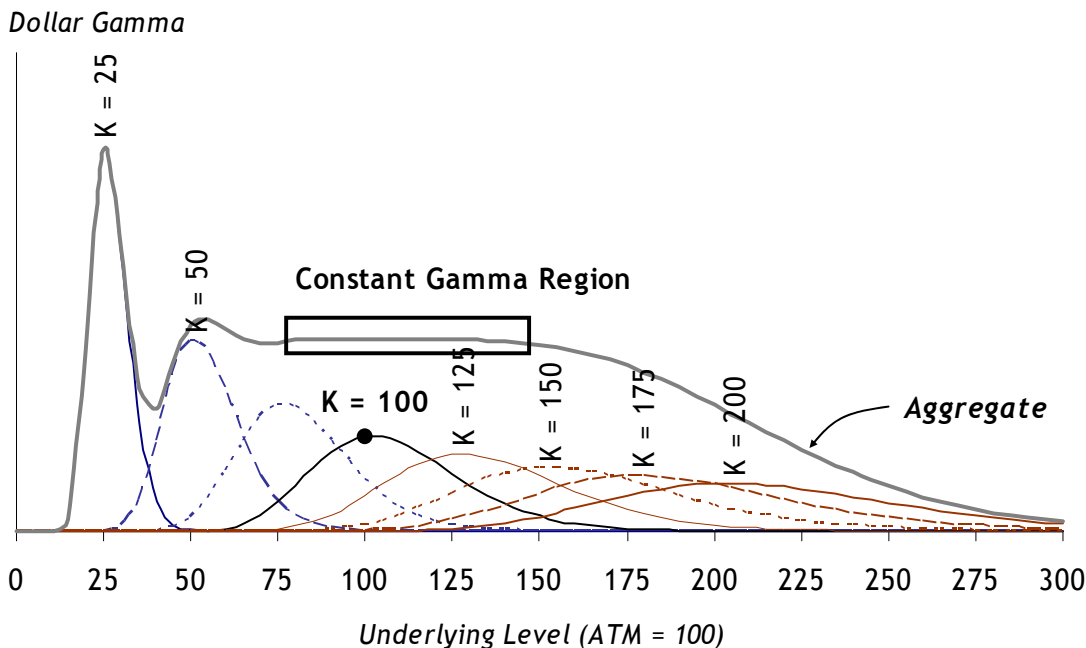
Interpretation

One might wonder what it means to create a derivative whose dollar gamma is constant. Dollar gamma is the standard gamma times S^2:

$$\Gamma^{\$}(S) = \frac{\partial^2 f}{\partial S^2} \times S^2$$

where f, S are the prices of the derivative and underlying, respectively. Thus, a constant dollar gamma means that for some constant a:

$$\frac{\partial^2 f}{\partial S^2} = \frac{a}{S^2}$$

The solution to this second-order differential equation is:

$$f(S) = -a\ln(S) + bS + c$$

where a, b, c are constants, and ln(.) the natural logarithm. In other words, the perfect static hedge for a variance swap would be a combination of the log-asset (a derivative which pays off the log-price of the underlying stock), the underlying stock and cash.

2.3. Valuation

Because a variance swap can be statically replicated with a portfolio of vanilla options, no particular modeling assumption is needed to determine its fair market value. The only model choice resides in the computation of the vanilla option prices — a task which merely requires a reasonable model of the implied volatility surface.

Assuming that one has computed the prices $p_0(k)$ and $c_0(k)$ of N^{puts} out-of-the money puts and N^{calls} out-of-the-money calls respectively, a quick proxy for the fair value of a variance swap of maturity T is given as:

JPMorgan

$$VarSwap_0 \approx \frac{2}{T}\left[\sum_{i=1}^{N^{puts}}\frac{p_0(k_i^{put})}{(k_i^{put})^2}\left(k_i^{put}-k_{i-1}^{put}\right)+\sum_{i=1}^{N^{calls}}\frac{c_0(k_i^{call})}{(k_i^{call})^2}\left(k_i^{call}-k_{i-1}^{call}\right)\right]$$
$$-PV_0(T)\times(K^{VS})^2$$

where $VarSwap_0$ is the fair present value of the variance swap for a variance notional of 1, K^{VS} is the strike, $PV_0(T)$ is the present value of \$1 at time T, k_i^{put} and k_i^{call} are the respective strikes of the i-th put and i-th call in percentage of the underlying *forward* price, with the convention $k_0 = 0$.

In the typical case where the strikes are chosen to be spaced equally apart, say every 5% steps, the expression between brackets is the sum of the put and call prices, weighted by the inverse of the squared strike, times the 5% step. Exhibit 2.3.1 below illustrates this calculation; in this example, the fair strike is around 16.62%, when a more accurate algorithm gave 16.54%. We also see that the fair strike is close to the 90% implied volatility (17.3%), as mentioned in Section 1.1.

Exhibit 2.3.1 — Calculation of the fair value of a variance swap through a replicating portfolio of puts and calls

In this example, the total hedge cost of the replicating portfolio is 2.7014% (=2/T * $\Sigma_i(w_i p_i)$), or 270.14 variance points. For a variance notional of 10,000, this means that the floating leg of the variance swap is worth €2,701,397.53. For a strike of 16.625 volatility points, and a 1-year present value factor of 0.977368853, the fixed leg is worth €2,701,355.88. Thus, the variance swap has a value close to 0.

Weight = $\frac{5\%}{Strike\%^2}$	Under-lying	Call / Put	Forward	Strike	Strike (%Forward)	Maturity	Implied Volatility	Price (%Notional)
20.00%	SX5E	P	2,935.02	1,467.51	50%	1Y	27.6%	0.04%
16.53%	SX5E	P	2,935.02	1,614.26	55%	1Y	26.4%	0.08%
13.89%	SX5E	P	2,935.02	1,761.01	60%	1Y	25.2%	0.15%
11.83%	SX5E	P	2,935.02	1,907.76	65%	1Y	24.0%	0.27%
10.20%	SX5E	P	2,935.02	2,054.51	70%	1Y	22.7%	0.46%
8.89%	SX5E	P	2,935.02	2,201.26	75%	1Y	21.4%	0.75%
7.81%	SX5E	P	2,935.02	2,348.01	80%	1Y	20.0%	1.17%
6.92%	SX5E	P	2,935.02	2,494.76	85%	1Y	18.7%	1.79%
6.17%	SX5E	P	2,935.02	2,641.51	90%	1Y	17.3%	2.67%
5.54%	SX5E	P	2,935.02	2,788.26	95%	1Y	16.0%	3.94%
2.50%	SX5E	P	2,935.02	2,935.02	100%	1Y	14.8%	5.74%
2.50%	SX5E	C	2,935.02	2,935.02	100%	1Y	14.8%	5.74%
4.54%	SX5E	C	2,935.02	3,081.77	105%	1Y	13.7%	3.37%
4.13%	SX5E	C	2,935.02	3,228.52	110%	1Y	12.9%	1.76%
3.78%	SX5E	C	2,935.02	3,375.27	115%	1Y	12.2%	0.81%
3.47%	SX5E	C	2,935.02	3,522.02	120%	1Y	11.9%	0.35%
3.20%	SX5E	C	2,935.02	3,668.77	125%	1Y	11.8%	0.15%
2.96%	SX5E	C	2,935.02	3,815.52	130%	1Y	11.9%	0.06%
2.74%	SX5E	C	2,935.02	3,962.27	135%	1Y	12.1%	0.03%
2.55%	SX5E	C	2,935.02	4,109.02	140%	1Y	12.5%	0.02%
2.38%	SX5E	C	2,935.02	4,255.77	145%	1Y	12.9%	0.01%
2.22%	SX5E	C	2,935.02	4,402.52	150%	1Y	13.4%	0.01%

Source: JPMorgan.

JPMorgan

3. Theoretical Insights

3.1. Idealized Definition of Variance

An idealized definition of annualized realized variance $W_{0,T}$ is given by:

$$W_{0,T} = \frac{1}{T}\big[\ln S, \ln S\big]_T$$

where S denotes the price process of the underlying asset and $[\ln S, \ln S]$ denotes the quadratic variation of $\ln S$. This definition is idealized in the sense that we implicitly assume that it is possible to monitor realized variance on a continuous basis. It can be shown that the discrete definition of realized variance given in Section 1.1 converges to the idealized definition above when moving to continuous monitoring.

This definition applies in particular to the classic Ito process for stock prices:

$$\frac{dS_t}{S_t} = \mu(t, S_t, \ldots)dt + \sigma(t, S_t, \ldots)dW_t$$

where the drift μ and the volatility σ are either deterministic or stochastic. In this case, the idealized definition of variance becomes:

$$W_{0,T} = \frac{1}{T}\int_0^T \sigma^2(t, S_t, \ldots)dt .$$

However, in the presence of jumps, the integral above only represents the *continuous contribution* to total variance, often denoted $\big[\ln S, \ln S\big]_T^c$. More details on the impact of jumps can be found in Section 3.4.

3.2. Hedging Strategies & Pricing

For ease of exposure, we assume in this section that dividends are zero and that the underlying price process S is a diffusion process. Moreover, let us assume that rates are deterministic. Let us introduce some notation: By S, we denote the non-discounted spot price process and by $\hat{S} = \frac{S}{B}$ we denote the discounted spot price process, where B refers to the deterministic money market account. It is important to note that $\big[\ln S, \ln S\big] = \big[\ln \hat{S}, \ln \hat{S}\big]$ when rates are deterministic. Moreover, the continuity of $\hat{S}$ together with Ito's formula yields:

$$\ln \hat{S}_t = \int_0^t \frac{1}{\hat{S}_u}d\hat{S}_u - \frac{1}{2}\big[\ln \hat{S}, \ln \hat{S}\big]_t \text{ for all } 0 \leq t \leq T.$$

Define for all $0 \leq t \leq T$:

$$\pi_t = \frac{1}{2}\big[\ln \hat{S}, \ln \hat{S}\big]_t + \ln \hat{S}_t = \int_0^t \frac{1}{\hat{S}_u}d\hat{S}_u .$$

We now explain how π_T, which is closely related to the payoff of a variance swap, can be replicated by continuous trading of the underlying and cash according to a self-financing strategy (V_0, φ, ψ), where V_0 is the initial value of the strategy, φ_t and ψ_t the quantities to be held in the underlying and cash at time t. The strategy is said to be self-financing because its mark-to-market value $V_t = V_0 + \varphi_t S_t + \psi_t B_t$ verifies:

JPMorgan

$$dV_t = \varphi_t dS_t + \psi_t dB_t$$

(In other words the change in value of the strategy between times t and $t + dt$ is computed as a mark-to-market P&L: change in asset price multiplied by the quantity held at time t. There is no addition or withdrawal of wealth.)

Self-financing strategy

One can verify that the following choice for (V_0, φ, ψ) is self-financing:

$$\begin{cases} V_0 = 0 \\ \varphi_t = \dfrac{1}{B_T \hat{S}_t} \\ \psi_t = \displaystyle\int_0^t \dfrac{1}{B_T \hat{S}_u} d\hat{S}_u - \dfrac{1}{B_T} \end{cases}$$

Let us point out a few important things:

— The self-financing strategy only replicates the terminal payoff π_T but it does not replicate π_t for $t < T$. It is indeed easy to see that $\pi_T = V_T$:

$$V_T = V_0 + \varphi_T S_T + \psi_T B_T = \frac{1}{B_T \hat{S}_T} S_T + \left(\int_0^T \frac{1}{B_T \hat{S}_u} d\hat{S}_u - \frac{1}{B_T} \right) B_T = \int_0^T \frac{1}{\hat{S}_u} d\hat{S}_u = \pi_T$$

However $\pi_t > V_t$ for t < T:

$$V_t = \frac{1}{B_T \hat{S}_t} S_t + \left(\int_0^T \frac{1}{B_T \hat{S}_u} d\hat{S}_u - \frac{1}{B_T} \right) B_t = \frac{B_t}{B_T} \int_0^T \frac{1}{\hat{S}_u} d\hat{S}_u = \frac{B_t}{B_T} \pi_t < \pi_t$$

— For the self-financing strategy to be predictable (i.e. for φ_t, ψ_t to be entirely determined based solely on the information available before time t), the assumption that rates are deterministic is crucial.

Pricing

Having identified a self-financing strategy we can proceed to price a variance swap by taking the risk-neutral expectation of π_T / B_T:

$$E\left[\frac{\pi_T}{B_T} \right] = E\left[\frac{1}{2B_T} [\ln \hat{S}, \ln \hat{S}]_T + \frac{1}{B_T} \ln \hat{S}_T \right] = E\left[\int_0^T \frac{1}{B_T \hat{S}_u} d\hat{S}_u \right] = 0$$

since $\hat{S}$ is assumed to be martingale under the risk-neutral measure. Whence:

$$E\left[\frac{1}{B_T} W_{0,T} \right] = -\frac{2}{T} E\left[\frac{1}{B_T} \ln \hat{S}_T \right]$$

At this point, it should be noted that this representation is valid only as long as we assume that the underlying stock price process is continuous and rates are deterministic. As soon as

JPMorgan

we deviate from this assumption, additional adjustments have to be made. For further details in this regard, see Sections 3.3 and 3.4.

Representation as a sum of puts and calls

In the previous paragraphs we showed that the annualized realized variance can be replicated with a static position in a log contract on the discounted stock price. However in general it is not possible to trade log contracts. Thus we need to obtain an alternative representation for the price of the variance swap using standard put and call options.

For this purpose, note that a twice differentiable payoff $f(S)$ can be re-written as follows:

$$f(S_T) = f(F_T) + f'(F_T)\left[(S_T - F_T)^+ - (F_T - S_T)^+\right]$$
$$+ \int_0^{F_T} f''(y)(y - S_T)^+ dy + \int_{F_T}^{+\infty} f''(y)(S_T - y)^+ dy$$

Here S_T denotes the spot price of the underlying and F_T denotes the forward price[10]. For details we refer to the Appendix in Carr-Madan (2002). Choosing $f(y) = \ln(y)$ and taking expectations yields:

$$E\left(\frac{1}{B_T} \ln \frac{S_T}{F_T}\right) = -\left[\int_0^1 \frac{1}{y^2} Put(y) dy + \int_1^{+\infty} \frac{1}{y^2} Call(y) dy\right]$$

where y now denotes forward moneyness, and $Put(y)$ or $Call(y)$ the price of a vanilla put or call expiring at time T. Whence:

$$E\left[\frac{1}{B_T} W_{0,T}\right] = \frac{2}{T}\left[\int_0^1 \frac{1}{y^2} Put(y) dy + \int_1^{+\infty} \frac{1}{y^2} Call(y) dy\right]$$

The interpretation of this formula is as follows: In case the stock price process S is a diffusion process, the annualized realized variance can be replicated by an infinite sum of static positions in puts and calls. Clearly, perfect replication is not possible since options for all strikes are not available. A more accurate representation would thus be a discretized version of the above (see Section 2.3 for an example.)

3.3. Impact of Dividends

When a stock pays a dividend, arbitrage considerations show that its price should drop by the dividend amount. This phenomenon results in a higher variance when the stock price is not adjusted for dividends, which is most often the case.

From a modeling standpoint, there are three standard ways to approach dividends: continuous dividend yield, discrete dividend yield, and discrete dollar dividend. In the following paragraphs we only focus on the first two cases:

— For continuous dividend yield, we consider the price process:

$$dS_t = (r_t - q_t)S_t dt + \sigma_t S_t dW_t$$

where r is a deterministic interest rate, q is a deterministic dividend yield, σ is either deterministic or stochastic.

---

[10] Since we assume zero dividends in this section, we have $F_T = S_0 B_T$.

JPMorgan

–   For discrete dividend yield, we consider the price process:

$$dS_t = r_t S_t dt + \sigma_t S_t dW_t - d\left(\sum_{t_j \le t} d_j S_{t_j} \right)$$

where r is a deterministic interest rate, σ is either deterministic or stochastic, and $d_1, \ldots, d_M$ are M discrete continously compounded dividend yields[11] paid at dates $t_1, \ldots, t_M$.

Continuous Monitoring

■ A continuous dividend yield has no impact on variance when monitoring is continuous. In this regard, observe that:

$$W_{0,T} = \frac{1}{T}\left[\ln S, \ln S\right]_T = \frac{1}{T}\left[\ln \hat{S}, \ln \hat{S}\right]_T$$

where $\hat{S} = \frac{S}{F}$ is the spot price normalized by the forward price. This is because the dividend yield q is assumed to be deterministic. Hence, there is clearly no impact due to continuous dividends[12]. The hedging strategy also remains the same.

■ Next, let us consider the impact of discrete dividends. In this case the stock price process S follows:

$$dS_t = r_t S_t dt + \sigma_t S_t dW_t - d\left(\sum_{t_j \le t} d_j S_{t_j} \right)$$

We now have:

$$W_{0,T} = \frac{1}{T}\left[\ln S, \ln S\right]_T = \frac{1}{T}\left[\ln \hat{S}, \ln \hat{S}\right]_T^c + \frac{1}{T}\sum_{t_j \le T} d_j^2$$

Let us have a closer look at the hedging strategy in the context of discrete dividend yields. For this purpose, define the total return process $G_t = S_t \exp\left(\sum_{t_j \le t} d_j\right)$ where dividends are reinvested in the stock. The discounted total return process $\hat{G} = G/B$ being a martingale we can use a similar hedging strategy as in Section 3.2 where the stock price process S is now replaced by G.

Discrete Monitoring

Consider a set of sampling dates $0 = t_0 < t_1 < \cdots < t_N = T$. For simplicity of presentation, we assume that the time intervals $\Delta t_i = t_i - t_{i-1}$ are all constant and equal to Δt. Recall the discrete definition of annualized variance without mean:

$$Variance = \frac{1}{T}\sum_{i=1}^{N}\left(\ln \frac{S_{t_i}}{S_{t_{i-1}}} \right)^2$$

---

[11] Note that we consider here gross yields rather than annualized yields in the discrete dividend case.

- Contrary to the continuous monitoring case, a continuous dividend yield has an impact on variance when monitoring is discrete. Consider the log return between t_{i-1} and t_i:

$$\ln \frac{S(t_i)}{S(t_{i-1})} = \left[r - q - \frac{1}{2}\sigma^2 \right] \Delta t + \sigma z \sqrt{\Delta t}$$

where $z \sim N(0,1), r = r_{t_{i-1}}, q = q_{t_{i-1}}, \sigma = \sigma_{t_{i-1}}$. Squaring the above yields:

$$\ln^2 \frac{S(t_i)}{S(t_{i-1})} = \left[r - q - \frac{1}{2}\sigma^2 \right]^2 \Delta^2 t + \sigma^2 z^2 \Delta t + 2\sigma z \left[r - q - \frac{1}{2}\sigma^2 \right] \Delta^{3/2} t$$

Because the expectation of z is nil and its variance $E(z^2)$ is one, we obtain:

$$E\left[\ln^2 \frac{S(t_i)}{S(t_{i-1})} \right] = \left[r - q - \frac{1}{2}\sigma^2 \right]^2 \Delta^2 t + \sigma^2 \Delta t$$

The relative impact of discrete monitoring on variance is thus:

$$\frac{E\left[\ln^2 \frac{S(t_i)}{S(t_{i-1})} \right]}{\sigma^2 \Delta t} = \frac{\left(r - q - \frac{1}{2}\sigma^2 \right)\Delta t}{\sigma^2}$$

At this point, it should be noted that even in the case where interest rates and dividends are assumed to be zero, we obtain some drift contribution in case of discretization. This is due to the σ^2 term in the numerator. Moreover, for $\Delta t \to 0$, the above expression implies that there is no contribution due to interest rates and continuous dividend yields — as already pointed out in the continuous monitoring case.

- We now specialize our considerations to the case of discrete dividends. Assuming that a discrete dividend d_j is paid between times t_{i-1} and t_i and carrying out similar calculations as in the previous paragraph yields the following expression for the expectation of the log return:

$$E\left[\ln^2 \frac{S_{t_i}}{S_{t_{i-1}}} \right] = \left(r - \frac{1}{\Delta t}d_j - \frac{1}{2}\sigma^2 \right)^2 \Delta^2 t + \sigma^2 \Delta t$$

As can be seen from this equation, the contribution of discrete dividends does not converge to zero for $\Delta t \to 0$. We also obtain that the relative contribution of the interest rate and the continuous dividend yield within a time interval Δt amounts to:

$$\frac{\left(r - \frac{1}{\Delta t}d_j - \frac{1}{2}\sigma^2 \right)^2 \Delta t}{\sigma^2}$$

---

[12] Note that this statement is true within a deterministic or stochastic volatility framework. In other frameworks (such as local volatility) volatility may depend on S and would thus be impacted by dividends.

3.4. Impact of Jumps

The purpose of this section is to analyze the impact of jumps, i.e. we no longer assume that the stock price process S follows a diffusion process and instead consider a jump diffusion process. For ease of exposure, we ignore interest rates and dividends:

$$\frac{dS_t}{S_{t-}} = \mu_t dt + \sigma_t dW_t + d\left(\sum_{n=1}^{N_t}(Y_n - 1)\right)$$

or:

$$d\ln S_t = \left(\mu_t - \frac{1}{2}\sigma_t^2\right)dt + \sigma_t dW_t + d\left(\sum_{n=1}^{N_t}\ln(Y_n - 1)\right)$$

where W, N and Y are independent. W is a standard Brownian motion, N is a Poisson process with intensity λ and (Y_n) are independent, identically distributed log-normal variables:

$$\begin{cases} dW_t \sim N(0, dt) \\ dN_t \sim P(\lambda dt) \\ Y_n = (1+k)e^{\delta G_n - \frac{1}{2}\delta^2}, \ G_n \sim N(0,1) \end{cases}$$

Parameters k, λ, δ can be interpreted as follows: k is the average jump size, λ controls the frequency of jumps, and δ is the jump size uncertainty (standard deviation.) Furthermore, the drift term μ_t is chosen such that $\hat{S} = S$ is a martingale, i.e.: $\mu_t = -\lambda k$. We then have for the annualized realized variance:

$$W_{0,T} = \frac{1}{T}[\ln S, \ln S]_T = \frac{1}{T}[\ln S, \ln S]_T^C + \frac{1}{T}\sum_{n=1}^{N_T}\ln^2(Y_n - 1)$$

And the expected variance under the risk-neutral measure becomes:

$$E[W_{0,T}] = \frac{1}{T}E\left[W_{0,T}^C\right] + \lambda\left[\left(\ln(1+k) - \frac{1}{2}\delta^2\right)^2 + \delta^2\right]$$

JPMorgan

Appendix A — A Review of Historical and Implied Volatility

Historical Volatility

The volatility of a financial asset (e.g. a stock) is the level of its price uncertainty, and is commonly measured by the standard deviation of its returns. For historical daily returns r_1, r_2, …, r_n, an estimate is given as:

$$\sigma_{Historical} = \sqrt{\frac{252}{n-1}\sum_{t=1}^{n}(r_t - \bar{r})^2}$$

where $\bar{r} = \frac{1}{n}\sum_{t=1}^{n}r_t$ is the mean return, and 252 is an annualization factor corresponding to the typical number of trading days in a year.

Historical volatility is also called realized volatility in the context of option trading and variance swaps.

Here it is assumed that the returns were independent and drawn according to the same random 'law' or distribution — in other words, stock prices are believed to follow a 'random walk.' In this case, the estimate is shown to be unbiased with vanishing error as the number of daily observations n increases.

The daily returns are typically computed in logarithmic terms in the context of options to remain consistent with Black-Scholes:

$$r_t = \ln\frac{P_t}{P_{t-1}}$$

where P_t is the price of the asset observed on day t, and ln(.) is the natural logarithm.

Implied volatility

Vanilla options on a stock are worth more when volatility is higher. Contrary to a common belief, this is not because the option has 'more chances of being in-the-money', but because the stock has more chances of being *higher* in-the-money, as illustrated in Exhibit A1.

In a Black-Scholes world, volatility is the only parameter which is left to the appreciation of the option trader. All the other parameters: strike, maturity, interest rate, forward value, are determined by the contract specifications and the interest rate and futures markets.

Thus, there is a one-to-one correspondence between an option's price and the Black-Scholes volatility parameter. Implied volatility is the value of the parameter for which the Black-Scholes theoretical price matches the market price, as illustrated in Exhibit A2.

Because of put-call parity, European calls and puts with identical characteristics (underlying, strike, maturity) must have the same implied volatility. This makes the distinction between volatilities implied from call or put prices irrelevant. In the case of American options, however, put-call parity does not always hold, and the distinction might be relevant.

For each strike and maturity there is a different implied volatility which can be interpreted as the market's expectation of future volatility between today and the maturity date *in the scenario implied by the strike*. For instance, out-of-the money puts are natural hedges against a market dislocation (such as caused by the 9/11 attacks on the World Trade Center) which entail a spike in volatility; the implied volatility of out-of-the money puts is thus higher

JPMorgan

than in-the-money puts. This phenomenon is known as volatility skew, as though the market expectations of uncertainty were *skewed* towards the downside.

An example of a volatility surface is given in Exhibit A3.
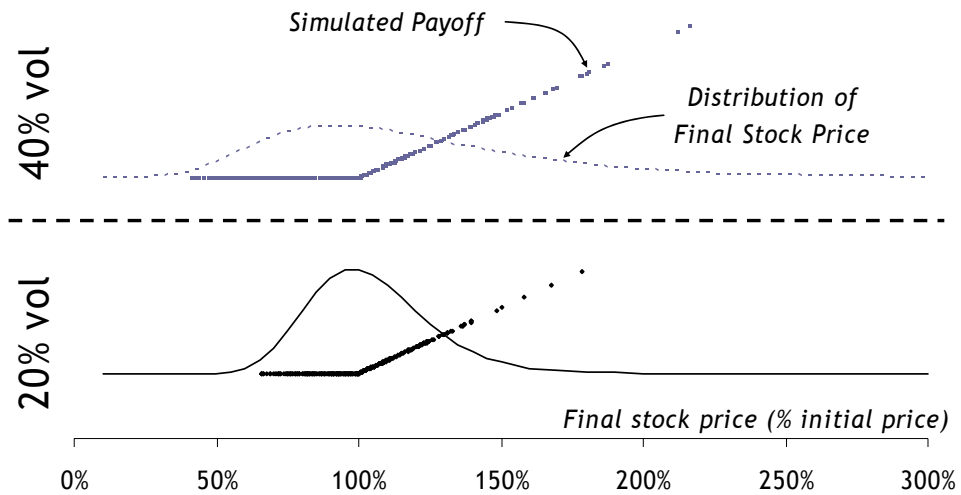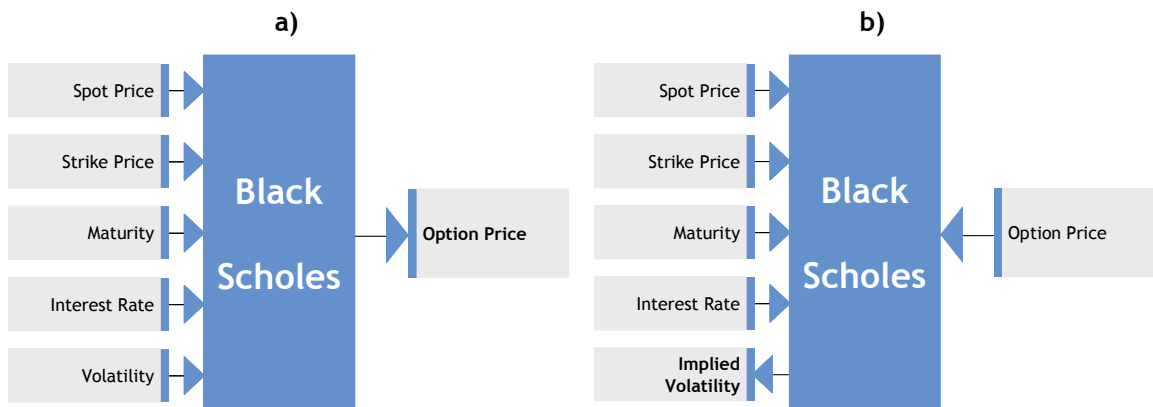
JUST WHAT YOU NEED TO KNOW ABOUT VARIANCE SWAPS
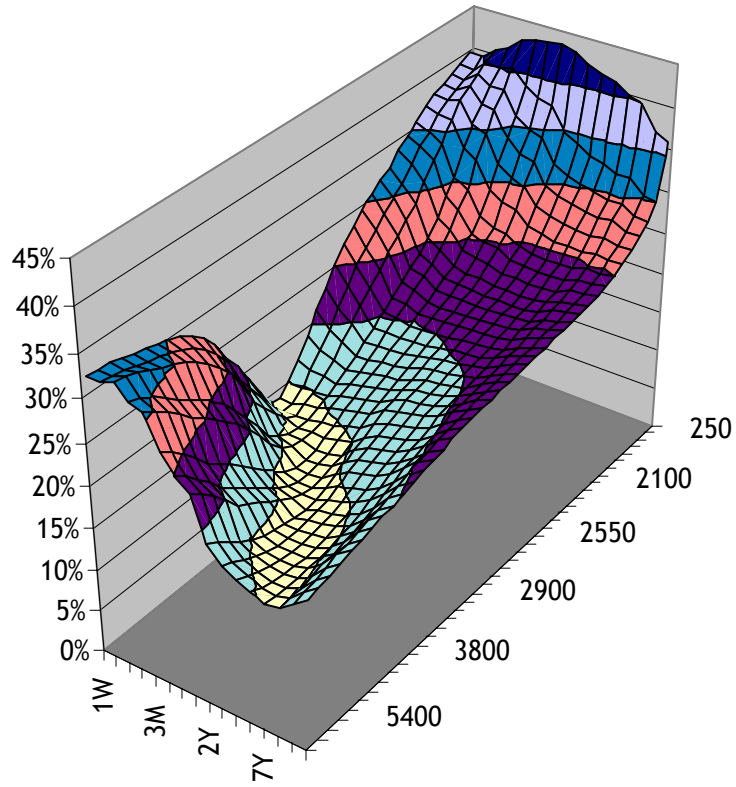
Source: JPMorgan.

JUST WHAT YOU NEED TO KNOW ABOUT VARIANCE SWAPS

JPMorgan

Appendix B — Relationship between Theta and Gamma

An intuitive approach

Consider the reduced P&L equation (Eq. 2) from Section 2.1

$$\text{Daily P\&L} = \frac{1}{2}\Gamma \times (\Delta S)^2 \quad + \quad \Theta \times (\Delta t) \qquad \textbf{(Eq. 2)}$$

In a fair game, the *expected* daily P&L is nil. This leaves us with:

$$\Theta \times \Delta t = -\frac{1}{2}\Gamma \times \mathbf{E}\!\left[(\Delta S)^2\right]$$

where **E**[.] denotes mathematical expectation[13]. Writing $(\Delta S)^2 = \left(\frac{\Delta S}{S}\right)^2 \times S^2$ yields:

$$\Theta \times \Delta t = -\frac{1}{2}\Gamma S^2 \times \mathbf{E}\!\left[\left(\frac{\Delta S}{S}\right)^2\right] \qquad \textbf{(Eq. B1)}$$

The quantity $\left(\frac{\Delta S}{S}\right)^2$ is the squared daily return on the underlying stock; taking expectation gives the stock variance over one day: $\sigma^2 \times \Delta t$. (Remember that implied volatility σ is given on an annual basis.) Replacing the expected squared return by its expression and dividing both sides of Equation B1 by Δt finally yields:

$$\Theta = -\frac{1}{2}\Gamma \sigma^2 S^2 .$$

By the books

Consider the Black-Scholes-Merton partial differential equation:

$$rf = \frac{\partial f}{\partial t} + rS\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} \qquad \textbf{(Eq. B2)}$$

where $f(t, S)$ is the value of the derivative at time t when the stock price is S, and r is the short-term interest rate.

Equation (B2) holds for all derivatives of the same underlying stock, and by linearity of differentiation any portoflio Π of such derivatives. Identifying the Greek letter corresponding to each partial derivative, we can rewrite Equation B2 as:

$$r\Pi = \Theta + rS\Delta + \frac{1}{2}\Gamma \sigma^2 S^2$$

In the case of a delta-hedged portfolio, we have $\Delta = 0$, whence:

$$\Theta = r\Pi - \frac{1}{2}\Gamma \sigma^2 S^2$$

Because the short-term rate is typically of the order of a few percentage points, the first term on the right-hand side is often negligible, and we have the approximate relationship:

$$\Theta \approx -\frac{1}{2}\Gamma \sigma^2 S^2 .$$

---

[13] Here we actually deal with conditional expectation upon the 'information' available at a certain point in time.
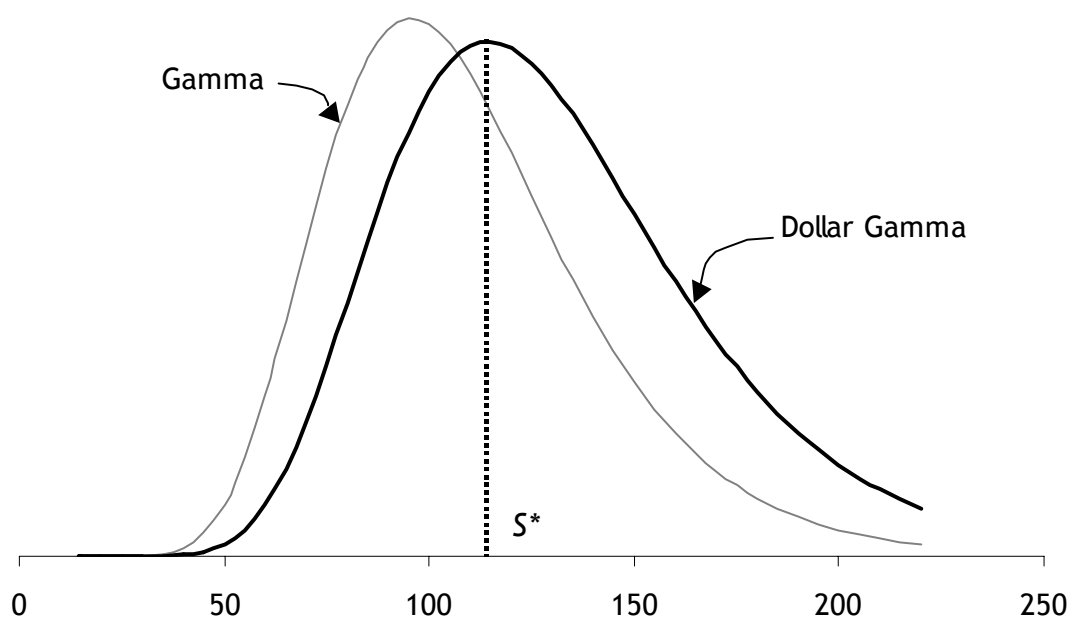
JPMorgan

Appendix C — Peak Dollar Gamma

When the interest rate is zero, the dollar gamma of a vanilla option with strike K, maturity T and implied volatility σ is given in function of the underlying level S as:

$$\Gamma^\$(S,K) = \frac{S}{\sigma\sqrt{2\pi T}} \exp\left(-\frac{(\ln(S/K)+0.5\sigma^2 T)^2}{2\sigma^2 T}\right)$$

In Exhibit C1 below we can see that the dollar gamma has a bell-shaped curve which peaks slightly after the 100 strike. It can indeed be shown that the peak is reached when S is equal to:

$$S^* = Ke^{\sigma\sqrt{T}-\sigma^2 T/2}$$

References & Bibliography

Allen, Harris (2001), *Volatility Vehicles*, JPMorgan Equity Derivatives Strategy Product Note.

Demeterfi, Derman, Kamal, Zou (1999), *More Than You Ever Wanted To Know About Volatility Swaps*, Goldman Sachs Quantitative Strategies Research Notes.

Carr, Madan (1998), 'Towards a Theory of Volatility Trading' in *VOLATILITY*, R.A. Jarrow Risk Books.

Gatheral (2002), *Case Studies in Financial Modelling Fall 2002*, NYU Courant Institute.

Hull (2000), *Options, Futures & Other Derivatives* 4th edition, Prentice Hall.

Musiela, Rutkowski (1997), *Martingale Methods in Financial Modelling*, Springer.

Vaillant (2001), *A Beginner's Guide to Credit Derivatives*, Working Paper, Nomura International.

Wilmott (1998), 'Discrete Hedging' in *The Theory and Practice of Financial Engineering*, Wiley.