


Box 11

VARIANCE SWAPS

In a financial market context, volatility is a measure of the extent of asset price fluctuations. It is a necessary input for various models used for pricing options and other financial instruments, and can be measured in many different ways. While so-called realised volatility measures the extent of past price variation, the volatility implied in options prices is used to gauge the market view of expected future price fluctuations. As the structures of financial instruments have

Figure B11.1 Variance swap cash flows

(variance swap strike = K^2 ; realised volatility = σ)

become increasingly complex, market participants need instruments which allow them to trade volatility in order to hedge structured transactions or to take purely directional views on volatility. While “plain vanilla” options contracts on underlying assets provide exposure to the volatility of the underlying asset, they are impure for hedging or taking positions on volatility because they simultaneously provide exposure to the direction of the underlying asset. Although hedging options according to the Black-Scholes prescription can remove the exposure to the underlying asset, so-called delta hedging is at best inaccurate because many of the

Black-Scholes assumptions are violated in practice. For instance, volatility cannot be accurately estimated, financial assets cannot be traded continuously, transaction costs cannot be ignored, markets sometimes move discontinuously, and liquidity is often a problem.

One response to such needs and challenges is the variance swap.¹ This instrument has gained a (perhaps unwarranted) reputation of exposing its investors to large risks with the potential to cause sizeable losses and exacerbate sharp market moves. However, while such risks cannot be excluded or underestimated, the appropriate use of variance swaps, like many other financial derivatives, can be of great benefit to sophisticated investors and market-makers. This Box describes the basic features of this instrument and outlines its possible uses for different types of market participant.

A variance swap is a forward contract on the difference between the variance delivery price, fixed at the inception of the contract, and the realised variance over the period of the swap.² Its structure is very similar to other swap contracts, whereby the counterparties to the trade agree at the time of entering into the contract on the fixed variance level for the contract period (usually the prevailing market implied variance, so that the swap’s market value at inception is zero). At maturity, the realised variance over the period is determined, and the difference, multiplied by the contract notional, is settled in cash. The variance swap buyer, holding a long position in variance (volatility), receives the payment from the variance swap seller (i.e. makes a profit) if the realised variance over the period is higher than the implied variance at the inception. If the realised variance over the period is lower than the implied variance at inception, the variance swap buyer must make the payment to the seller (who holds a short position in variance/volatility) and thus realises a loss.³ The variance swap payout profile is asymmetrical, as the long position gains more when volatility is rising than it loses when volatility declines by the same amount. In other words, the variance swap is convex with respect to volatility. Because realised volatility cannot be less than zero, a long variance swap position has a known maximum loss.⁴ The maximum loss on a short variance

1 See K. Demeterfi, E. Derman, M. Kamal and J. Zou (1999), “More Than You Ever Wanted to Know about Volatility Swaps (But Less Than Can Be Said)”, Goldman Sachs, *Quantitative Research Strategies Notes*, March.

2 Variance is a square of volatility. While volatility swaps are also quoted and traded, it is less onerous to replicate and hedge variance in practice. This is why the market for variance swaps has developed to a greater extent than volatility swaps.

3 Realised volatility is systematically lower than implied volatility. This is due to the fact that implied volatility levels include a risk premium for tail events when there are unexpectedly large volatility spikes.

4 This turns out to be $\frac{1}{2}$ times the strike price times the vega notional, where the vega is a coefficient measuring the sensitivity of an option value to a change in volatility.

swap is theoretically unlimited. However, because realised volatility can reach very high levels in case of market upheaval, potential losses are often limited by the inclusion of a cap on volatility.

Although variance swaps were first traded as early as the late 1990s, liquid markets for these instruments did not fully take off for some time because of a lack of a universally accepted pricing methodology. Only after robust pricing models had been introduced did the market develop properly.⁵ Initially, variance swaps were offered on the most liquid equity indices, such as the S&P 500, the EURO STOXX 50, the DAX and the FTSE, and indices still remain the most common underlying assets in variance swaps. Variance swaps on individual stocks, especially the more liquid constituents of the popular equity indices, are also traded, and even though the market for equity indices as the underlying is the most advanced, there is no obstacle, at least in theory, to variance swaps being traded on other asset classes, including foreign exchange, commodities or interest rates.

The variance swap market has grown steadily in recent years, with institutional investors increasingly using variance swaps for hedging purposes or for portfolio diversification. Numerous opportunities exist for variance swaps to be used in trading or hedging strategies, thus increasing the choices available to market participants to express their market views or to hedge their exposures. For example, life assurance companies now offer many products with guaranteed benefits (e.g. variable annuities, with-profit funds), and these expose them to short volatility positions that may be offset by using variance swaps. In addition, both outright directional volatility trades as well as spread trades exploiting relative value in volatilities across different assets or time periods are possible through the use of variance swaps. Because increases in volatility can persist for a period of time after a sharp (especially downward) movement in prices, equity investors can buy variance swaps to offset the risk of a fall in the value of their holdings if the market declines. Variance swaps can also be used for hedging purposes by market-makers wanting to dispose of their exposures from various client transactions.

Some concerns have been expressed about the potential that trading in variance swaps could amplify market volatility and perhaps even create adverse market dynamics. For instance, if a market-maker in the variance swaps market needs to hedge exposures arising from trading a variance swap with another market-maker or a client and therefore uses a large portfolio of options with dynamic delta hedging to replicate the variance swap's payout, this may result in different price dynamics compared to the "usual" delta hedging arising from a simple options trade. Because realised variance is determined on the basis of closing prices for each day of the contract period, variance swap market-makers who delta hedge their positions only need to hedge their exposure against the closing prices. If the underlying asset experiences a large daily move, this hedging action has the potential to amplify asset price changes still further, thus increasing volatility. On top of this, expectations of such hedging activities may prompt other market participants to take the same positions earlier during a trading session so that they too can benefit from price changes near the market close. This "feedback" effect has been frequently mentioned as one of the factors that amplified equity market volatility in May 2006.

⁵ These pricing models are based on the theory that a variance swap payout can be replicated using a large number of plain vanilla options at various strikes, complemented by dynamic delta hedging in the underlying asset. The number of options at each strike level is inversely correlated to the strike level, i.e. at low strike levels the portfolio holds a large number of options, and vice versa. In practice, only selected strike levels are used because of the low liquidity of deep out-of-the money options and high transaction costs.