

CHAPTER 14

Tools for Volatility Engineering, Volatility Swaps, and Volatility Trading

1. Introduction

Liquid instruments that involve *pure* volatility trades are potentially very useful for market participants who have natural exposure to various volatilities in their balance sheet or trading book. The classical option strategies discussed in Chapter 10 have serious drawbacks in this respect. When a trader takes a position or hedges a risk, he or she expects that the random movements of the underlying would have a *known* effect on the position. The underlying may be random, but the *payoff function* of a well-defined contract or a position has to be known. Payoff functions of most *classical* volatility strategies are not invariant to underlying risks, and most volatility instruments turn out to be *imperfect* tools for isolating this risk. Even when traders' anticipations come true, the trader may realize that the underlying volatility payoff functions have changed due to movements in *other* variables. Hence, classical volatility strategies cannot provide satisfactory hedges for volatility exposures. The reason for this and possible solutions are the topics of this chapter.

Traditional volatility trades used to involve buying and selling portfolios of call and put options, straddles or strangles, and then possibly *delta*-hedging these positions. But such volatility positions were not *pure* and this led to a search for volatility tools whose payoff function would depend on the volatility risk *only*. This chapter examines two of the pure volatility instruments that were developed—variance and volatility swaps. They are interesting for at least two reasons: First, volatility is an important risk for market practitioners, and ways of hedging and pricing such risks have to be understood. Second, the discussion of volatility swaps constitutes a good example of the basic principles that need to be followed when devising new instruments.

The chapter uses *variance* swaps instead of volatility swaps to conduct the discussion. Although markets in general use the term *volatility*, it is more appropriate to think in terms of *variance*, the square of volatility. Variance is the second centered moment of a random variable,

and it falls more naturally from the formulas used in this chapter. For example, volatility (i.e., standard deviation) instruments require convexity adjustments, whereas variance instruments in general do not. Thus, when we talk about *vega*, for example, we refer to *variance vega*. This is the sensitivity of the option's price with respect to σ^2 , not σ . In fact, in the heuristic discussion in this chapter, the term volatility and variance are used interchangeably.

2. Volatility Positions

Volatility positions can be taken with the purpose of hedging a volatility exposure or speculating on the future behavior of volatility. These positions require instruments that isolate volatility risk as well as possible. To motivate the upcoming discussion, we introduce two examples that illustrate traditional volatility positions.

2.1. Trading Volatility Term Structure

We have seen several examples for strategies associated with shifts in the interest rate term structure. They were called *curve steepening* or *curve flattening* trades. It is clear that similar positions can be taken with respect to *volatility* term structures as well. Volatilities traded in markets come with different maturities. As with the interest rate term structure, we can buy one "maturity" and sell another "maturity," as the following example shows.¹

EXAMPLE:

[A] dealer said he was considering selling short-dated 25-delta euro puts/dollar calls and buying a longer-dated straddle. A three-month straddle financed by the sale of two 25-delta one-month puts would have cost 3.9% in premium yesterday.

These volatility plays are attractive because the short-dated volatility is sold for more than the cost of the longer-maturity options.

In this particular example, the anticipations of traders concern not the level of an asset price or return, but, instead, the volatility associated with the price. Volatility over one interval is bought using the funds generated by selling the volatility over a different interval.

Apparently, the traders think that *short-dated* euro volatility will fall *relative to* the long-dated euro volatility. The question is to what extent the positions taken will meet the traders' needs, *even when their anticipations are borne out*. We will see that the payoff function of this position is not invariant to changes in the underlying euro/dollar exchange rate.

2.2. Trading Volatility across Instruments

Our second example is from the interest rate sector and involves another "arbitrage" position on volatility. The trader buys the volatility of one risk and sells a related volatility on a different risk. This time, the volatilities in question do not belong to different time periods, but instead are generated by different *instruments*.

¹ The term "arbitrage" is used here in the sense financial markets use it and not in the sense of academic analysis. The following arbitrage positions may have zero cost and have a relatively high probability of succeeding, but the gains are by no means risk-free.

EXAMPLE:

Dealers are looking at the spreads between euro cap-floor straddle and swaption straddle volatility to take advantage of a 5% volatility difference in the 7-year area. Proprietary traders are selling a two-year cap-floor straddle starting in six years with vols close to 15%. The trade offers a good pick-up over the five-year swaption straddle with volatility 10%. This compares with a historical spread closer to 2%.

Cap-floors and swaptions are instruments on interest rates. There are both similarities and differences between them. We will study them in more detail in the next chapter. Selling a cap-floor straddle will basically be *short* interest rate volatility. In the example, the traders were able to take this position at 15% volatility. On the other hand, buying a swaption amounts to a *long* position on volatility. This was done at 10%, which gives a *volatility spread* of about 5%. The example states that the latter number has historically been around 2%. Hence, by selling the spread the traders would benefit from a future narrowing of differences between the volatilities of the two instruments.

This position's payoff is not invariant to interest rate trajectories. Even when volatilities behave as anticipated, the path followed by the level of interest rates may result in unexpected volatility.² The following discussion intends to clarify why such positions on volatility have serious weaknesses and require meticulous risk management. We will consider pure volatility positions later.

3. Invariance of Volatility Payoffs

In previous chapters, convexity was used to isolate *volatility* as a risk. In Chapters 8 and 9, we showed how to convert the volatility of an underlying into “cash,” and with that took the first steps toward volatility engineering.

Using the methods discussed in Chapters 8 and 9, a trader can hedge and risk-manage exposures with respect to volatility movements. Yet, these are positions influenced by variables *other* than volatility. Consider a speculative position taken by an *investor*.

Let S_t be a risk factor and suppose an investor *buys* S_t volatility at time t_0 for a future period denoted by $[t_1, T]$, T being the expiration of the contract. As in every long position, this means that the investor is anticipating an increase in *realized* volatility during this period. If realized volatility during $[t_1, T]$ exceeds the volatility “purchased” at time t_0 , the investor will benefit. Thus far this is not very different from other long positions. For example, a trader buys a stock and benefits if the stock price goes up. He or she can buy a fixed receiver swap and benefit if the swap value goes up (the swap rate goes down). Similarly, in our present case, we receive a “fixed” volatility and benefit if the actual volatility exceeds this level.

By buying call or put options, straddles, or strangles, and then *delta*-hedging these positions, the trader will, in general, end up with a long position that benefits if the realized volatility increases, as was shown in Chapters 8 and 9. Yet, there is one major difference between such volatility positions and positions taken on other instruments such as stocks, swaps, forward rate agreements (FRAs), and so on. Consider Figure 14-1a, that shows a long position on a stock funded by a money market loan. As the stock price increases, the position benefits by the amount $S_{t_1} - S_{t_0}$. This potential payoff is known and depends *only* on the level of S_{t_1} . In fact,

² We must point out that there are differences between cap-floor volatilities and swaption volatilities. In fact, this 4% spread may very well be due to these factors. Also, such positions become even more complicated with the existence of a volatility smile.

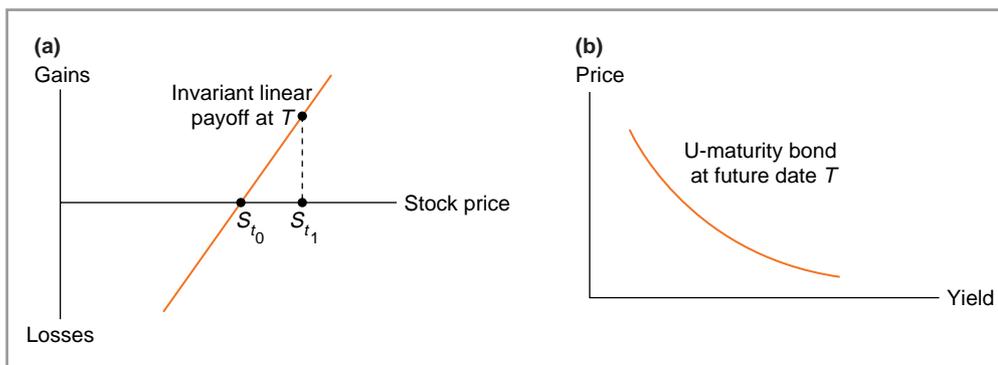


FIGURE 14-1

it depends on S_t linearly. In Figure 14-1b we have a *short*-dated discount bond position. As the yield decreases, the position gains. Again, we *know* how much the position will be making or losing, depending on the movements in the yield, y_t , if convexity gains are negligible.

A volatility position taken via, say, straddles, is fundamentally different from this as the payoff diagram will move depending on the path followed by variables *other* than volatility. For example, a change in (1) interest rates, (2) the underlying asset price, or (3) the level of *implied* volatility may lead to different payoffs at the same *realized* volatility level.

Variance (volatility) swaps, on the other hand, are pure volatility positions. Potential gains or losses in positions taken with these instruments depend *only* on what happens to realized volatility until expiration. This chapter shows how volatility engineering can be used to set up such contracts and to study their pricing and hedging. We begin with imperfect volatility positions.

3.1. Imperfect Volatility Positions

In financial markets, a *volatility position* is often interpreted to be a static position taken by buying and selling straddles, or a dynamically maintained position that uses straddles or options. As mentioned previously, these volatility positions are not the right way to price, hedge, or risk-manage volatility exposure. In this section, we go into the reasons for this. We consider a simple position that consists of a dynamically hedged single-call option.

3.1.1. A Dynamic Volatility Position

Consider a volatility exposure taken through a dynamically maintained position using a plain vanilla call. To simplify the exposition, we impose the assumptions of the Black-Scholes world where there are no dividends; the interest rate, r , and implied volatility, σ , are constant; there are no transaction costs; and the underlying asset follows a geometric process. Then the arbitrage-free value of a European call $C(S_t, t)$ will be given by the Black-Scholes formula:

$$C(S_t, t) = S_t \int_{-\infty}^{d_1} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} dx - e^{-r(T-t)} K \int_{-\infty}^{d_2} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} dx \quad (1)$$

where S_t is the spot price, and K is the strike. The $d_i, i = 1, 2$, are given by

$$d_i = \frac{\log \frac{S_t}{K} \pm \frac{1}{2}\sigma^2(T-t) + r(T-t)}{\sigma\sqrt{T-t}} \quad (2)$$

For simplicity, and without loss of generality, we let

$$r = 0 \quad (3)$$

This simplifies some expressions and makes the discussion easier to follow.³

Now consider the following simple experiment. A trader uses the Black-Scholes setting to take a dynamically hedged *long* position on implied volatility. Implied volatility goes up. Suppose the trader tracks the gains and losses of the position using the corresponding variance-*vega*. What would be this trader's possible gains in the following specific case? Consider the following simple setup.

1. The parameters of the position are as follows:

$$\text{Time to expiration} = 0.1 \quad (4)$$

$$K = S_{t_0} = 100 \quad (5)$$

$$\sigma = 20\% \quad (6)$$

Initially we let $t_0 = 0$.

2. The trader expects an increase in the *implied* volatility from 20% to 30%, and considers taking a *long* volatility position.
3. To *buy* into a volatility position, the trader borrows an amount equal to $100 C(S_t, t)$, and buys 100 calls at time t_0 with funding cost $r = 0$.⁴
4. Next, the position is *delta*-hedged by short-selling C_s units of the underlying per call to obtain the familiar exposure shown in Figure 14-2.

In this example, there are about 1.2 months to the expiration of this option, the option is at-the-money, and the initial implied volatility is 20%.

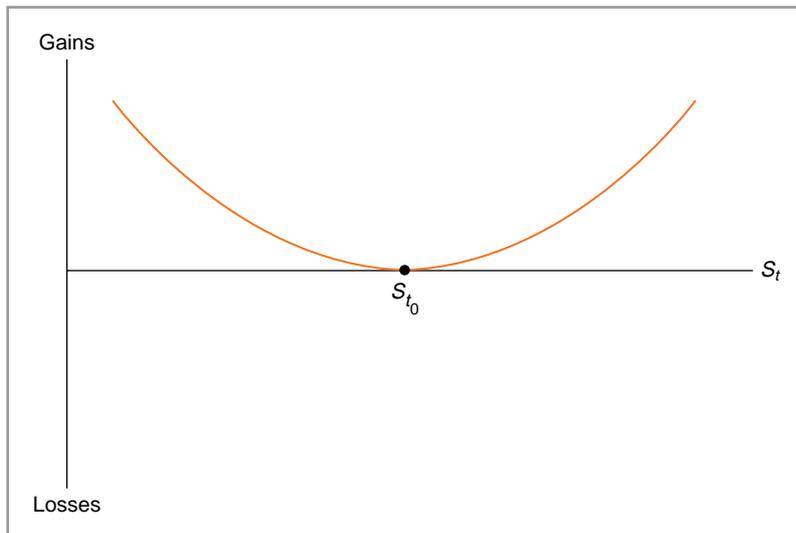


FIGURE 14-2

³ This is a useful assumption for discussing volatility trading.

⁴ Remember that an identical position could be taken by buying puts. We take calls simply as an example.

It turns out that in this environment, even when the trader's anticipations are borne out, the payoffs from the volatility position may vary significantly, depending on the path followed by S_t . The implied volatility may move from 20% to 30% as anticipated, but the position may not pay the expected amount. The following example displays the related calculations.

EXAMPLE:

We can calculate the relevant Greeks and payoff curves using Mathematica. First, we obtain the initial price of the call as

$$C(100, t_0) = 2.52 \quad (7)$$

Multiplying by 100, the total position is worth \$252. Then, we get the implied delta of this position by first calculating the S_t -derivative of $C(S_t, t)$ evaluated at $S_{t_0} = 100$, and then multiplying by 100:

$$100 \left(\frac{\partial C(S_t, t)}{\partial S_t} \right) = 51.2 \quad (8)$$

Hence, the position has +51-delta. To hedge this exposure, the trader needs to short 51 units of the underlying and make the net delta exposure approximately equal to zero.

Next, we obtain the associated gamma and the (variance) vega of the position at t_0 . Using the given data, we get

$$\text{Gamma} = 100 \left[\frac{\partial^2 C(S_t, t)}{\partial S_t^2} \right] = 6.3 \quad (9)$$

$$\text{Variance vega} = 100 \left[\frac{\partial C(S_t, t)}{\partial \sigma^2} \right] = 3,152 \quad (10)$$

The change in the option value, given a change in the (implied) variance, is given approximately by

$$100[\partial C(S_t, t)] \cong (3, 152)\partial\sigma^2 \quad (11)$$

This means that, everything else being constant, if the implied volatility rises suddenly from 20% to 30%, the instantaneous change in the option price will depend on the product of these numbers and is expected to be

$$\begin{aligned} 100[\partial C(S_t, t)] &\cong 3,152(.09 - .04) & (12) \\ &= 157.6 & (13) \end{aligned}$$

In other words, the position is expected to gain about \$158, if everything else remained constant.

The point is that the trader was long implied volatility, expecting that it would increase, and it did. So if the volatility *does* go up from 20% to 30%, is this trader guaranteed to gain the \$157.6? Not necessarily. Let us see why not.

Even in this simplified Black-Scholes world, the (variance) vega is a function of S_t, t, r , as well as σ^2 . Everything else is not constant and the S_t may follow any conceivable trajectory. But, and this is the important point, when S_t changes, this in turn will make the vega change as well. The following table shows the possible values for variance-vega depending on the value assumed by S_t , within this setting.⁵

⁵ The numbers in the table need to be multiplied by 100.

S_t	Vega
80	0.0558
90	7.4666
100	31.5234
110	10.6215
120	0.5415

Thus, if the expectations of the trader are fulfilled, the implied volatility increases to 30%, but, at the same time, if the underlying price moves away from the strike, say to $S_{t_1} = 80$, the same calculation will become approximately:

$$\text{Vega} (\partial\sigma^2) \cong 5.6(.09 - .04) \quad (14)$$

$$= 0.28 \quad (15)$$

Hence, instead of an anticipated gain of \$157.6, the trader could realize almost no gain at all. In fact, if there are costs to maintaining the volatility position, the trader may end up losing money. The reason is simple: as S_t changes, the option's sensitivity to implied volatility, namely the *vega*, changes as well. It is a function of S_t . As a result, the outcome is very different from what the trader was originally expecting.

For a more detailed view on how the position's sensitivity moves when S_t changes, consider Figure 14-3 where we plot the partial derivative:

$$100 \frac{\partial \text{Variance vega}}{\partial S_t} \quad (16)$$

Under the present conditions, we see that as long as S_t remains in the vicinity of the strike K , the trader has some exposure to volatility changes. But as soon as S_t leaves the neighborhood of K , this exposure drops sharply. The trader may think he or she has a (variance) volatility position,

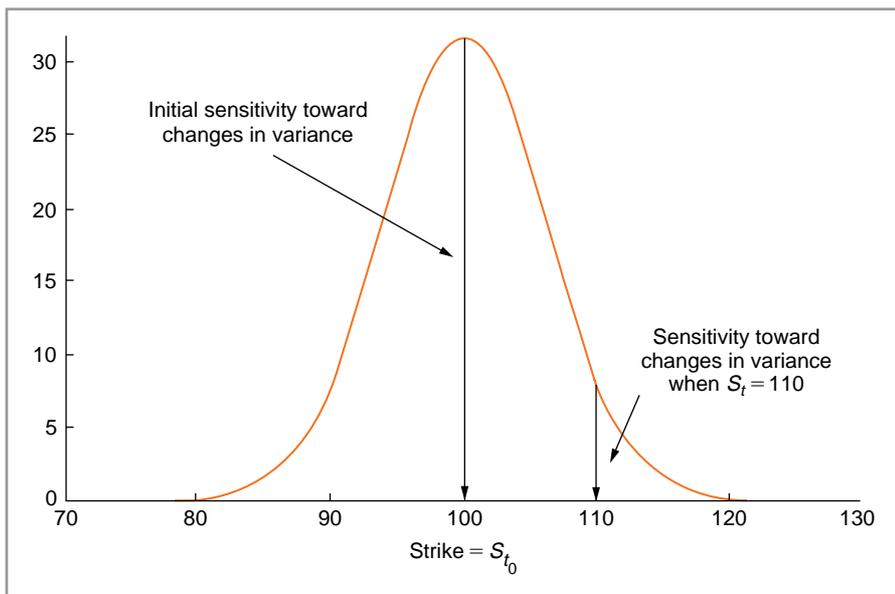


FIGURE 14-3

but, in fact, the position costs money, and may not have any significant variance exposure as the underlying changes right after the trade is put in place. Thus, such classical volatility positions are imperfect ways of putting on volatility trades or hedging volatility exposures.

3.2. Volatility Hedging

The outcome of such volatility positions may also be unsatisfactory if these positions are maintained as a hedge against a *constant* volatility exposure in another instrument. According to what was discussed, movements in S_t can make the hedge *disappear* almost completely and the trader may hold a naked volatility position in the end. An institution that has volatility exposure may use a hedge only to realize that the hedge may be *slipping* over time due to movements unrelated to volatility fluctuations.

Such slippage may occur for more reasons than just a change in S_t . In reality, there are also (1) smile effects, (2) interest rate effects, and (3) shifts in correlation parameters in some instruments. Changes in these can also cause the classical volatility payoffs to move away from initially perceived levels.

3.3. A Static Volatility Position

If a *dynamic delta-neutral* option position loses its exposure to movements in σ^2 and, hence, ceases to be useful as a hedge against volatility risk, do static positions fare better?

A classic position that has volatility exposure is buying (selling) ATM straddles. Using the same numbers as above, Figure 14-4 shows the joint payoff of an ATM call and an ATM put struck at $K = 100$. This position is made of two plain vanilla options and may suffer from a similar defect. The following example discusses this in more detail.

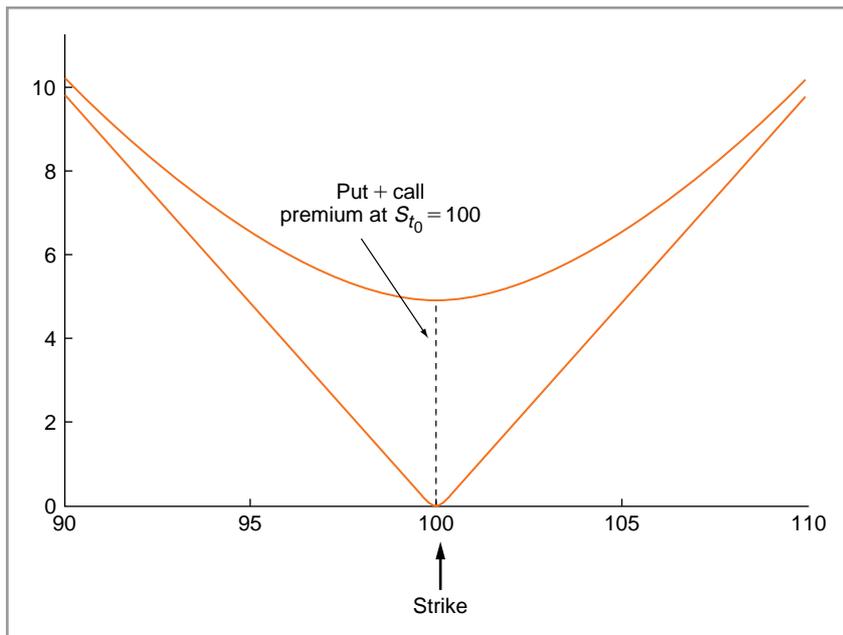


FIGURE 14-4

EXAMPLE:

As in the previous example, we choose the following numerical values:

$$S_{t_0} = 100, r = 0, T - t_0 = .1 \quad (17)$$

The initial volatility is 20%, which means that

$$\sigma^2 = .04 \quad (18)$$

We again look at the sensitivity of the position with respect to movements in some variables of interest. We calculate the variance vega of the portfolio:

$$V(S_t, t) = 100\{\text{ATM Put} + \text{ATM Call}\} \quad (19)$$

by taking the partial:

$$\text{Straddle vega} = 100 \frac{\partial V(S_t, t)}{\partial \sigma^2} \quad (20)$$

Then, we substitute the appropriate values of S_t, t, σ^2 in the formula. Doing this for some values of interest for S_t , we obtain the following sensitivity factors:

S_t	Vega
80	11
90	1493
100	6304
110	2124
120	108

According to these numbers, if S_t stays at 100 and the volatility moves from 20% to 30%, the static position's value increases approximately by

$$\partial \text{Straddle} \cong 6,304(.09 - .04) \quad (21)$$

$$= 315.2 \quad (22)$$

As expected, this return is about twice as big as in the previous example. The straddle has more sensitivity to volatility changes. But, the option's responsiveness to volatility movements is again not constant, and depends on factors that are external to what happens to volatility. The table shows that if S_t moves to 80, then even when the trader's expectation is justified and volatility moves from 20% to 30%, the position's mark-to-market gains will go down to about 0.56.

Figure 14-5 shows the behavior of the straddle's sensitivity with respect to implied volatility for different values of S_t . We see that the volatility position is again not invariant to changes in external variables. However, there is one major difference from the case of a dynamically maintained portfolio. Static non-*delta*-hedged positions using straddles will benefit from *actual* (realized) movements in S_t . For example, if the S_t stays at 80 until expiration date T , the put leg of the straddle would pay 20 and the static volatility position would gain. This is regardless of how the *vega* of the position changed due to movements in S_t over the interval $[t_0, T]$.

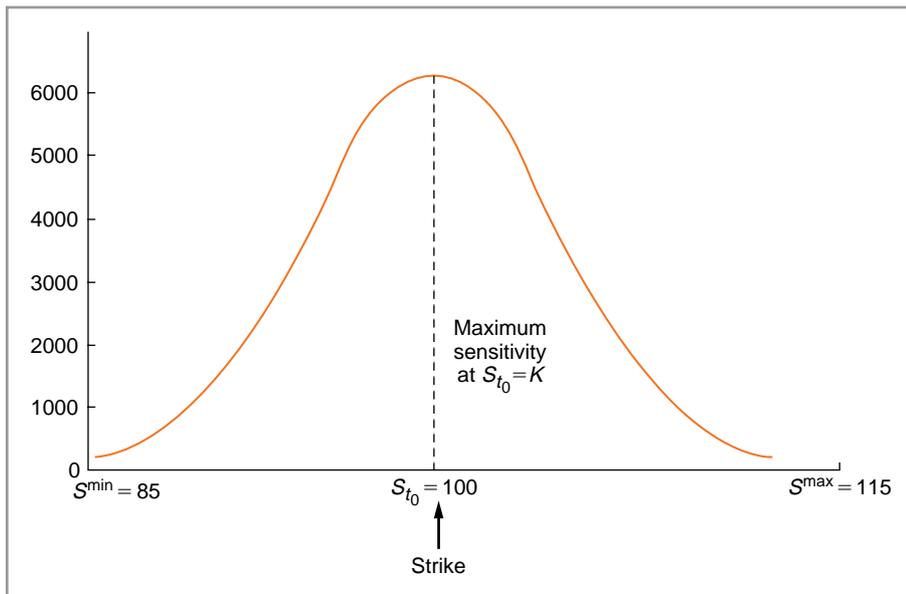


FIGURE 14-5

4. Pure Volatility Positions

The key to finding the right way to hedge volatility risk or to take positions in it is to isolate the “volatility” completely, using existing liquid instruments. In other words, we have to construct a *synthetic* such that the value of the synthetic changes *only* when “volatility” changes. This position should not be sensitive to variations in variables other than the underlying volatility. The exposure should be invariant. Then, we can use the synthetic to take volatility exposures or to hedge volatility risk. Such volatility instruments can be quite useful.

First, we know from Chapters 11 and 12 that by using options with different strikes we can essentially create *any* payoff that we like—if options with a broad range of strikes exist and if markets are complete. Thus, we should, in principle, be able to create pure volatility instruments by using judiciously selected option portfolios.

Second, if an option position’s *vega* drops suddenly once S_t moves away from the strike, then, by combining options of *different* strikes appropriately, we may be able to obtain a *portfolio* of options whose *vega* is more or less insensitive to movements in S_t . Heuristically speaking, we can put together small portions of smooth *curves* to get a desired *horizontal line*.

When we follow these steps, we can create pure volatility instruments. Consider the plot of the *vega* of three plain vanilla European call options, two of which are out-of-the-money. The options are identical in all respects, except for their strike. Figure 14-6 shows an example. Three σ^2 sensitivity factors for the strikes $K_0 = 100$, $K_1 = 110$, $K_2 = 120$ are plotted. Note that *each* variance *vega* is very sensitive to movements in S_t , as discussed earlier. Now, what happens when we consider the portfolio made of the *sum* of all three calls? The sensitivity of the portfolio,

$$V(S_t, t) = \{C(S_t, t, K_0) + C(S_t, t, K_1) + C(S_t, t, K_2)\} \quad (23)$$

again varies as S_t changes, but less. So, the direction taken is correct except that the previous portfolio did not optimally combine the three options. In fact, according to Figure 14-6, we should have combined the options by using different *weights* that depend on their respective

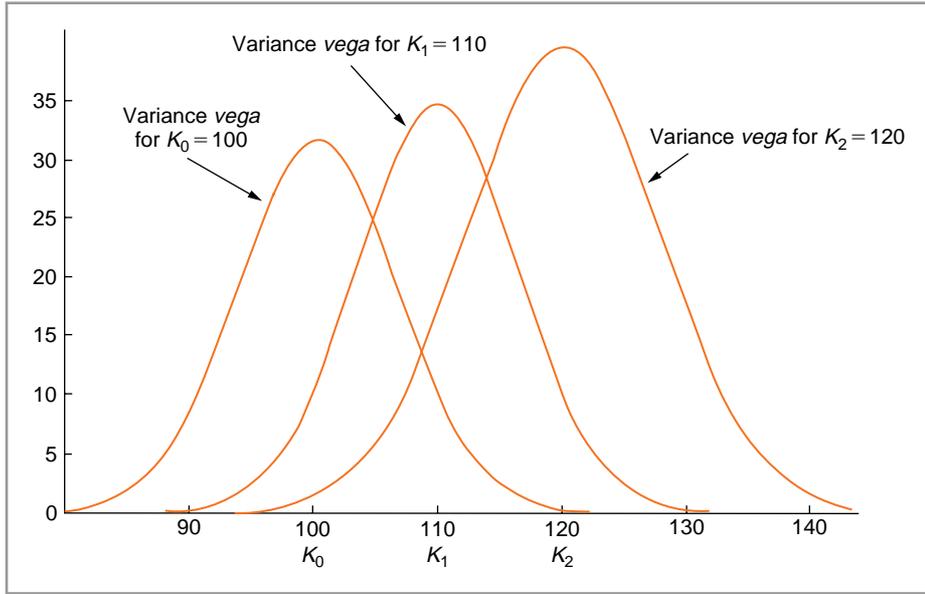


FIGURE 14-6

strike price. The more out-of-the-money the option is, the higher should be its weight, and the more it should be present in the portfolio.

Hence, consider the new portfolio where the weights are inversely proportional to the square of the strike K ,

$$V(S_t, t) = \frac{1}{K_0^2} C(S_t, t, K_0) + \frac{1}{K_1^2} C(S_t, t, K_1) + \frac{1}{K_2^2} C(S_t, t, K_2) \quad (24)$$

The variance *vega* of this portfolio that uses the parameter values given earlier, is plotted in Figure 14-7. Here, we consider a suitable $0 < \epsilon$, and the range

$$K_0 - \epsilon < S_t < K_2 + \epsilon \quad (25)$$

Figure 14-7 shows that the *vega* of the portfolio is approximately constant over this range when S_t changes. This suggests that more options with different strikes can be added to the portfolio, weighting them by the corresponding strike prices. In the example below we show these calculations.

EXAMPLE:

Consider the portfolio

$$V(S_t, t) = \left[\frac{1}{80^2} C(S_t, t, 80) + \frac{1}{90^2} C(S_t, t, 90) + \frac{1}{100^2} C(S_t, t, 100) \right. \quad (26)$$

$$\left. + \frac{1}{110^2} C(S_t, t, 110) + \frac{1}{120^2} C(S_t, t, 120) \right] \quad (27)$$

This portfolio has an approximately constant vega for the range

$$80 - \epsilon < S_t < 120 + \epsilon \quad (28)$$

By including additional options with different strikes in a similar fashion, we can lengthen this section further.

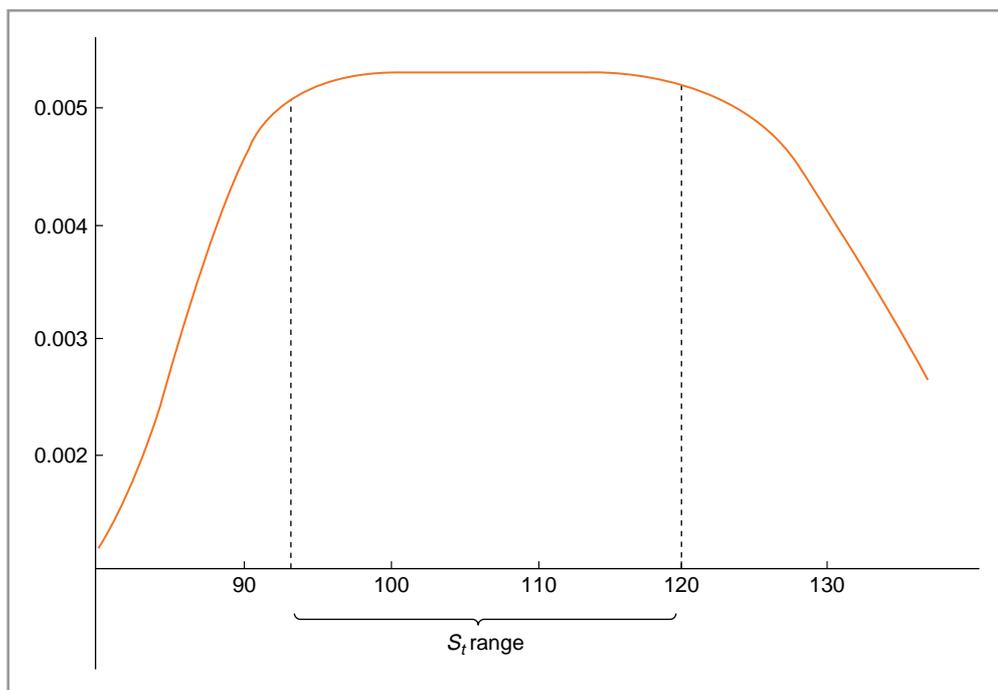


FIGURE 14-7

We have, in fact, found a way to create synthetics for volatility positions using a portfolio of liquid options with varying strikes, where the portfolio options are weighted by their respective strikes.

4.1. Practical Issues

In our attempt to obtain a pure volatility instrument, we have essentially followed the same strategy that we have been using all along. We constructed a *synthetic*. But this time, instead of matching the cash flows of an instrument, the synthetic had the purpose of matching a particular *sensitivity* factor. It was put together so as to have a constant (variance) *vega*.

Once a constant *vega* portfolio is found, the payoff of this portfolio can be expressed as an approximately linear function of σ^2

$$V(\sigma^2) = a_0 + a_1\sigma^2 + \text{small} \quad (29)$$

with

$$a_1 = \frac{\partial V(\sigma^2, t)}{\partial \sigma^2} \quad (30)$$

as long as S_t stays within the range

$$S^{\min} = K_0 < S_t < K_n = S^{\max} \quad (31)$$

Under these conditions, the volatility position will look like any other long (or short) position, with a positive slope a_1 .

The portfolio with a constant (variance) *vega* can be constructed using vanilla European calls and puts. The rules concerning synthetics discussed earlier apply here also. It is important

that elements of the synthetic be liquid; therefore liquid calls and puts have to be selected. The previous discussion referred only to calls. Practical applications of the procedure involve puts as well. This brings us to two somewhat complicated issues. The first has to do with the smile effect. The second concerns liquidity.

4.1.1. The Smile Effect

Suppose we form a portfolio at time t_0 that has a constant *vega* as long as S_t stays in a reasonable range

$$S^{\min} < S_t < S^{\max} \quad (32)$$

Under these conditions, the portfolio consists of options with different “moneyness” properties, and the volatility parameter in the option pricing formulas may depend on K if there is a volatility smile. In general, as K decreases, the implied $\sigma(K)$ would increase for constant S_t . Under these conditions, the trader needs to accurately determine the smile and the way to model it before the portfolio is formed.

4.1.2. Liquidity Problems

From the preceding it follows that we need to select out-of-the-money options for the synthetic since they are more liquid. But as time passes, the moneyness of these options changes and this affects their liquidity. Those options that become in-the-money are now less liquid. Other options that were not originally included in the synthetic become more liquid. Even though the replicating portfolio was static, the illiquidity of the constituent options may become a drawback in case the position needs to be unwound.

5. Volatility Swaps

One instrument that has invariant exposure to fluctuations in (realized) volatility is the *volatility swap*. In this section, we introduce this concept and in the next, we provide a simple framework for studying it.

A variance swap is, in many ways, just like any other swap. The parties exchange *floating* risk against a risk fixed at the contract origination. In this case, what is being swapped is not an interest rate or a return on an equity instrument, but the volatilities that correspond to various risk factors.

In the following section we move to a more technical discussion of volatility (variance) swaps. However, we emphasize again that the discussion will proceed using the variance rather than the volatility as the underlying.

5.1. A Framework for Volatility Swaps

Let S_t be the underlying price. The time- T_2 payoff $V(T_1, T_2)$ of a variance swap with a notional amount, N , is given by the following:

$$V(T_1, T_2) = [\sigma_{T_1, T_2}^2 - F_{t_0}^2] (T_2 - T_1)N \quad (33)$$

where σ_{T_1, T_2} is the *realized* volatility rate of S_t during the interval $t \in [T_1, T_2]$, with $t < T_1 < T_2$. It is similar to a “floating” rate, and will be observed only when time T_2 arrives. The F_{t_0} is the “fixed” S_t volatility *rate* that is quoted at time t_0 by markets. This has to be multiplied by $(T_2 - T_1)$ to get the appropriate volatility for the contract period. N is the notional amount that

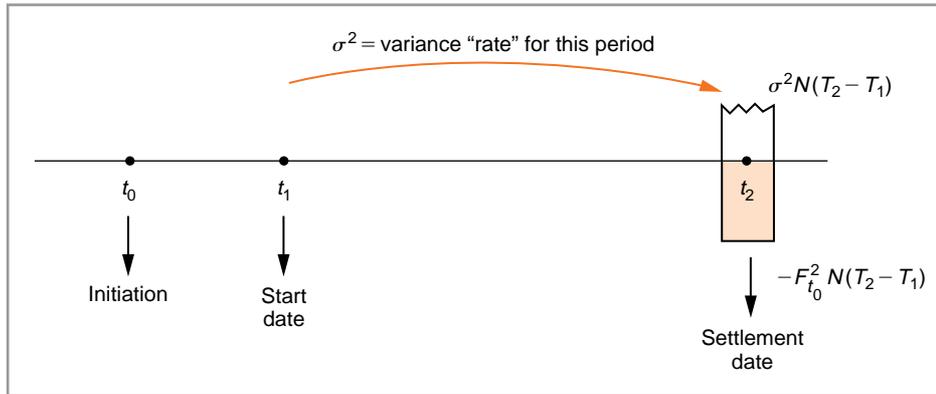


FIGURE 14-8

needs to be determined at contract initiation. At time t_0 , the $V(T_1, T_2)$ is unknown. The swap is set so that the time- t_0 “expected value” of the payoff, denoted by $V(t_0, T_1, T_2)$ is zero. At initiation, no cash changes hands:

$$V(t_0, T_1, T_2) = 0 \tag{34}$$

Thus, variance swaps are similar to a vanilla swap in that a “floating” $\sigma_{T_1, T_2}^2 (T_2 - T_1) N$ is received against a “fixed” $(T_2 - T_1) F_{t_0}^2 N$.

The cash flows implied by a variance swap are shown in Figure 14-8. The contract is initiated at time t_0 , and the start date is T_1 . It matures at T_2 . The “floating” volatility (variance) is the total volatility (variance) of S_t during the entire period $[T_1, T_2]$. F_{t_0} has the subscript t_0 , and, hence, has to be determined at time t_0 . We look at the two legs of the swap in more detail.

5.1.1. Floating Leg

Volatility positions need to be taken with respect to a well-defined time interval. After all, the volatility *rate* is like an interest rate: It is defined for specific time interval. Thus, we subdivide the period $[T_1, T_2]$ into equal subintervals, say, days:

$$T_1 = t_1 < t_2 \dots < t_n = T_2 \tag{35}$$

with

$$t_i - t_{i-1} = \delta \tag{36}$$

and then define the realized variance for period δ as

$$\sigma_{t_i}^2 \delta = \left[\frac{S_{t_i} - S_{t_{i-1}}}{S_{t_{i-1}}} - \mu \delta \right]^2 \tag{37}$$

where $i = 1, \dots, n$.⁶ Here, μ is the expected rate of change of S_t during a year. This parameter can be set equal to zero or any other estimated mean. Regardless of the value chosen, μ needs to

⁶ Of course, there are many other ways to define these “short-period” volatilities. Some of the recent research uses the estimated variance of daily price changes during a trading day, for example.

be carefully defined in the contract. If μ is zero, then the right-hand side is simply the squared returns during intervals of length δ .

Adding the marginal variances for successive intervals, σ_{T_1, T_2}^2 is equal to

$$(\sigma_{T_1, T_2}^2) (T_2 - T_1) = \sum_{i=1}^n \left[\frac{S_{t_i} - S_{t_{i-1}}}{S_{t_i}} - \mu\delta \right]^2 \quad (38)$$

Thus, σ_{T_1, T_2}^2 represents the realized percentage variance of the S_t during the interval $[T_1, T_2]$.

If the intervals become smaller and smaller, $\delta \rightarrow 0$, the last expression can be written as

$$(\sigma_{T_1, T_2}^2) (T_2 - T_1) = \int_{T_1}^{T_2} \left[\frac{1}{S_t} dS_t - \mu dt \right]^2 \quad (39)$$

$$= \int_{T_1}^{T_2} \sigma_t^2 dt \quad (40)$$

This formula defines the realized volatility (variance). It is a random variable at time t_0 , and can be viewed as the floating leg of the swap. Obviously, such floating volatilities can be defined for any interval in the future and can then be exchanged against a “fixed” leg.

5.1.2. Determining the Fixed Volatility

Determining the fixed volatility, F_{t_0} , will give the fair value of the variance swap at time t_0 . How do we obtain the numerical value of F_{t_0} ? We start by noting that the variance swap is *designed* so that its fair value at time t_0 is equal to zero. Accordingly, the $F_{t_0}^2$ is that number (variance), which makes the fair value of the swap equal zero. This is a basic principle used throughout the text and it applies here as well.

We use the fundamental theorem of asset pricing and try to find a proper arbitrage-free measure \tilde{P} such that

$$E_{t_0}^{\tilde{P}} [\sigma_{T_1, T_2}^2 - F_{t_0}^2] (T_2 - T_1) N = 0 \quad (41)$$

What could this measure \tilde{P} be? Suppose markets are complete.

We assume that the continuously compounded risk-free spot rate r is constant. The random process σ_{T_1, T_2}^2 is, then, a nonlinear function of S_u , $T_1 \leq u \leq T_2$, only:

$$\sigma_{T_1, T_2}^2 (T_2 - T_1) = \int_{T_1}^{T_2} \left[\frac{1}{S_t} dS_t - \mu dt \right]^2 \quad (42)$$

Under some conditions, we can use the normalization by the money market account and let \tilde{P} be the risk-neutral measure.⁷ Then, from equation (41), taking the expectation inside the brackets and arranging, we get

$$F_{t_0}^2 = E_{t_0}^{\tilde{P}} [\sigma_{T_1, T_2}^2] \quad (43)$$

This leads to the pricing formula

$$F_{t_0}^2 = \frac{1}{T_2 - T_1} E_{t_0}^{\tilde{P}} \left[\int_{T_1}^{T_2} \left[\frac{1}{S_t} dS_t - \mu dt \right]^2 \right] \quad (44)$$

⁷ We remind the reader that this contract will be settled at time T_2 .

Therefore, to determine $F_{t_0}^2$ we need to evaluate the expectation under the measure \tilde{P} of the integral of σ_t^2 . The discrete time equivalent of this is given by

$$F_{t_0}^2 = \frac{1}{T_2 - T_1} E_{t_0}^{\tilde{P}} \left[\sum_{i=1}^n \left[\frac{S_{t_i} - S_{t_{i-1}}}{S_{t_{i-1}}} - \mu \delta \right]^2 \right] \quad (45)$$

Given a proper arbitrage-free measure, it is not difficult to evaluate this expression. One can use Monte Carlo or tree methods to do this once the arbitrage-free dynamics is specified.

5.2. A Replicating Portfolio

The representation using the risk-neutral measure can be used for pricing. But how would we hedge a variance swap? To create the right hedge, we need to find a replicating portfolio. We discuss this issue using an alternative setup. This alternative has the side advantage of the financial engineering interpretation of some mathematical tools being clearly displayed. The following model starts with Black-Scholes assumptions.

The trick in hedging the variance swap lies in isolating σ_{T_1, T_2}^2 in terms of observable (traded) quantities. This can be done by obtaining a proper synthetic. Assume a diffusion process for S_t :

$$dS_t = \mu(S_t, t)S_t dt + \sigma(S_t, t)S_t dW_t \quad t \in [0, \infty) \quad (46)$$

where W_t is a Wiener process defined under the probability \tilde{P} . The diffusion parameter $\sigma(S_t, t)$ is called local volatility. Now consider the nonlinear transformation:

$$Z_t = f(S_t) = \log(S_t) \quad (47)$$

We apply Ito's Lemma to set up the dynamics (i.e., the SDE) for this new process Z_t :

$$dZ_t = \frac{\partial f(S_t)}{\partial S_t} dS_t + \frac{1}{2} \frac{\partial^2 f(S_t)}{\partial S_t^2} \sigma(S_t, t)^2 S_t^2 dt \quad t \in [0, \infty) \quad (48)$$

which gives

$$d \log(S_t) = \frac{1}{S_t} \mu(S_t, t) S_t dt - \frac{1}{2 S_t^2} \sigma(S_t, t)^2 S_t^2 dt + \sigma(S_t, t) dW_t \quad t \in [0, \infty) \quad (49)$$

where the S_t^2 term cancels out on the right-hand side. Collecting terms, we obtain

$$d \log(S_t) = \left[\mu(S_t, t) - \frac{1}{2} \sigma(S_t, t)^2 \right] dt + \sigma(S_t, t) dW_t \quad (50)$$

Notice an interesting result. The dynamics for dS_t/S_t and $d \log(S_t)$ are almost the same *except* for the factor involving $\sigma(S_t, t)^2 dt$. This means that we can subtract the two equations from each other and obtain

$$\frac{dS_t}{S_t} - d \log(S_t) = \frac{1}{2} \sigma(S_t, t)^2 dt \quad t \in [0, \infty) \quad (51)$$

This operation has *isolated* the instantaneous percentage local volatility on the right-hand side. But, what we need for the variance swap is the *integral* of this term. Integrating both sides we get

$$\int_{T_1}^{T_2} \left[\frac{1}{S_t} dS_t - d \log(S_t) \right] = \frac{1}{2} \int_{T_1}^{T_2} \sigma(S_t, t)^2 dt \quad (52)$$

We now take the integral on the left-hand side,

$$\int_{T_1}^{T_2} d \log(S_t) = \log(S_{T_2}) - \log(S_{T_1}) \quad (53)$$

We use this and rearrange to obtain the result:

$$2 \left[\int_{T_1}^{T_2} \frac{1}{S_t} dS_t \right] - 2 \log \left(\frac{S_{T_2}}{S_{T_1}} \right) = \int_{T_1}^{T_2} \sigma(S_t, t)^2 dt \quad (54)$$

We have succeeded in isolating the percentage total variance for the period $[T_1, T_2]$ on the right-hand side. Given that S_t is an asset that trades, the expression on the left-hand side replicates this variance.

5.3. The Hedge

The interpretation of the left-hand side in equation (54) is quite interesting. It will ultimately provide a hedge for the variance swap. In fact, the integral in the expression is a good example of what Ito integrals often mean in modern finance. Consider

$$\int_{T_1}^{T_2} \frac{1}{S_t} dS_t \quad (55)$$

How do we interpret this expression?

Suppose we would like to maintain a long position that is made of $\frac{1}{S_t}$ units of S_t held during each infinitesimally short interval of size dt , and for all t . In other words, we purchase $\frac{1}{S_t}$ units of the underlying at time t and hold them during an infinitesimal interval dt . Given that at time t , S_t is observed, this position can easily be taken. For example, if $S_t = 100$, we can buy 0.01 units of S_t at a total cost of 1 dollar. Then, as time passes, S_t will change by dS_t and the position will gain or lose dS_t dollars for every unit purchased. We readjust the portfolio, since the S_{t+dt} will presumably be different, and the portfolio needs to be $\frac{1}{S_{t+dt}}$ units long.

The resulting gains or losses of such portfolios during an infinitesimally small interval dt are given by the expression⁸

$$\frac{1}{S_t} (S_{t+dt} - S_t) = \frac{1}{S_t} dS_t \quad (56)$$

Proceeding in a similar fashion for all subsequent intervals dt , over the entire period $[T_1, T_2]$, the gains and losses of such a dynamically maintained portfolio add up to

$$\int_{T_1}^{T_2} \frac{1}{S_t} dS_t \quad (57)$$

The integral, therefore, represents the *trading gains or losses* of a dynamically maintained portfolio.⁹

⁸ The use of dt here is heuristic.

⁹ In fact, this interpretation can be generalized quite a bit. Often the stochastic integrals in finance have a structure such as

$$\int_{T_1}^{T_2} f(S_t) dS_t$$

These can be interpreted as trading gains or losses of *dynamically* maintaining $f(S_t)$ units of the asset that have price S_t .

The second integral on the left-hand side of equation (52)

$$\int_{T_1}^{T_2} d \log(S_t) = \log(S_{T_2}) - \log(S_{T_1}) \quad (58)$$

is taken with respect to time t , and is a standard integral. It can be interpreted as a *static* position. In this case, the integral is the payoff of a contract written at time T_1 , which pays, at time T_2 , the difference between the unknown $\log(S_{T_2})$ and the known $\log(S_{T_1})$. This is known as a *log contract*. The long and short positions in this contract are logarithmic functions of S_t .

In a sense, the left-hand side of equation (54) provides a hedge of the variance contract. If the trader is short the variance swap, he or she would also maintain a dynamically adjusted long position on S_t and be short a static log contract. This assumes complete markets.

6. Some Uses of the Contract

The variance (volatility) swaps are clearly useful for taking positions with volatility exposure and hedging. But, each time a new market is born, there are usually further developments beyond the immediate uses. We briefly mention some further applications of the notions developed in this chapter.

First of all, the F_t^2 , which is the fixed leg of the variance swap, can be used as a benchmark in creating new products. It is important to realize, however, that this price was obtained using the risk-neutral measure and that is *not* necessarily an unbiased forecast of future volatility (variance) for the period $[T_1, T_2]$. Just like the FRA market prices, the F_t will include a risk premium. Still, it is the proper price on which to write volatility options.

The pricing of the variance swap does not necessarily give a volatility that will equal the implied volatility for the same period. Implied volatility comes with a smile and this may introduce another wedge between F_t and the ATM volatility.

Finally, the F_t^2 should be a good indicator for risk-managing volatility exposures and also options books.

The following reading illustrates the development of this market.

EXAMPLE:

A striking illustration of the increasing awareness of volatility among the hedge fund community is the birth of pure volatility funds. But just as notable as the introduction of specialist volatility investment vehicles is the growing realization among regular directional hedge funds of the need to manage their volatility positions.

“As people become aware of volatility, they are increasingly looking to hedge or trade the vega,” said a participant from a directional hedge fund.

Convertible arbitrage funds have also been getting in on the act as they come to fully understand the concept of vega. Volatility is a major factor in the pricing of convertible bonds.

Investment banks have responded to an increased hedge fund interest in volatility by providing new straightforward volatility structures.

The best example of the new breed of simple volatility products is the volatility swap. These are cash-settled forward bets on market volatility which allow the investor to set up a pure volatility trade with a dealer. When the customer sells volatility, the dealer agrees to pay a fixed volatility rate on a notional amount for a certain period. In return,

the investor agrees to pay the annualized realized volatility for the S&P500 for the life of the swap.

At maturity, the two income streams are netted and the counterparties exchange the difference in whichever direction is appropriate. This type of product encourages hedge fund volatility activity because it offers them a simpler method of trading vega.

Normal volatility trades, such as caps and floors, leave investors exposed to underlying price risk. As the market moves towards the strike price, the gamma effect in hedging the position may cause the investor to lose more on the hedging than he makes on the volatility rate. Careful book management is necessary to control this risk. Most directional hedge funds have so many things to look at that they haven't always got the time, inclination or understanding to trade volatility using the traditional products. "Volatility swaps turn vega into something that people can easily grasp and manage," said one directional hedge fund commentator. (*IFR*, December 31, 1998)

Volatility trading, volatility hedging, and arbitraging all fall within a sector that is still in the process of development. In the next chapter we will see some new difficulties and new positions associated with them.

7. Which Volatility?

This chapter dealt with *four* notions of volatility. These must be summarized and distinguished clearly before we move on the discussion of the volatility smile in the next chapter.

When market professionals use the term "volatility," chances are they refer to Black-Schole's *implied* volatility. Otherwise, they will use terms such as *realized* or *historical* volatility. *Local volatility* and *variance swap volatility* are also part of the jargon. Finally, *cap-floor volatility* and *swaption volatility* are standard terms in financial markets.

Implied volatility is simply the value of σ that one would plug into the Black-Scholes formula to obtain the fair market value of a plain vanilla option as observed in the markets. For this reason, it is more correct to call it *Black-Scholes implied vol* or *Black volatility* in the case of interest rate derivatives. It is quite conceivable for a professional to use a different formula to price options, and the volatility implied by this formula would naturally be different. The term implied volatility is, thus, a formula-dependent variable.

We can attach the following definitions to the term "volatility."

- First, there is the class of *realized volatilities*. This is closest to what is contained in statistics courses. In this case, there is an observed or to-be-observed data set, a "sample," $\{x_1, \dots, x_n\}$, which can be regarded as a realization of a possibly vector-stochastic process, x_t , defined under some real-world probability P . The process x_t has a second moment

$$\sigma_t = \sqrt{E_t^P [(x_t - E_t^P[x_t])^2]} \quad (59)$$

We can devise an estimator to estimate this σ_t . For example, we can let

$$\hat{\sigma}_t = \sqrt{\frac{\sum_{i=0}^m (x_{t-i} - \bar{x}_t^m)^2}{m}} \quad (60)$$

where \bar{x}_t^m is the m -period sample mean:

$$\bar{x}_t^m = \frac{\sum_{i=0}^m x_{t-i}}{m} \quad (61)$$

Such volatilities measure the actual real-world fluctuations in asset prices or risk factors. One example of the use of this volatility concept was shown in this chapter. The σ_t^2 defined earlier represented the floating leg of the variance swap discussed here.

- The next class is *implied volatility*.¹⁰ There is an observed market price. The market practitioner has a pricing formula (e.g., Black-Scholes) or procedure (e.g., implied trees) for this price. Then, implied volatility is that “volatility” number, or series of numbers, which must be plugged into the formula in order to recover the fair market price. Thus, let $F(S_t, t, r, \sigma_t, T)$ be the Black-Scholes price for a European option written on the underlying S_t , with interest rates r and expiration T . At time t , σ_t represents the implied volatility if we solve the following equation (nonlinearly) for σ_t :

$$F(S_t, t, r, \sigma_t, T) = \text{Observed price} \quad (62)$$

This implied volatility may differ from the realized volatility significantly, since it incorporates any *adjustments* that the trader feels he or she should make to expected realized volatility. Implied volatility may be systematically different than realized volatility if volatility is *stochastic* and if a *risk premium* needs to be added to volatility quotes. Violations of Black-Scholes assumptions may also cause such a divergence.

- *Local volatility* is used to represent the function $\sigma(\cdot)$ in a stochastic differential equation:

$$dS(t) = \mu(S, t)dt + \sigma(S, t)S_t dW_t \quad t \in [0, \infty) \quad (63)$$

However, local volatility has a more specific meaning. Suppose options on S_t trade in all strikes, K , and expirations T , and that the associated arbitrage-free prices, $\{C(S_t, t, K, T)\}$, are observed for all K, T . Then the function $\sigma(S_t, t)$ is the local volatility, if the corresponding SDE successfully replicates all these observed prices either through a Monte Carlo or PDE pricing method.

In other words, *local volatility* is a concept associated with calibration exercises. It can be regarded as a generalization of Black-Scholes implied volatility. The implied volatility replicates a *single* observed price through the Black-Scholes formula. The local volatility, on the other hand, replicates an entire *surface* of options indexed by K and T , through a pricing method. As a result, we get a *volatility surface* indexed by K and T , instead of a single number as in the case of Black-Scholes implied volatility.

- Finally, in this chapter we encountered the *variance swap volatility*. This referred to the expectation of the average future squared deviations. But, because the expectation used the risk-neutral measure, it is different from real-world volatility.

Discussions of the volatility smile relate to these volatility notions. The implied volatility is obviously of interest to most traders but it cannot exist independently of realized volatility. It is natural to expect a close relationship between the two concepts. Also, as volatility trading develops, more and more instruments are written that use the realized volatility as some kind of underlying risk factor for creating new products. The variance swap was only one example.

8. Conclusions

This chapter provided a brief introduction to a sector that may, in the future, play an even more significant role in financial market strategies. Our purpose was to show how we can isolate the

¹⁰ This definition could be a little misleading since these days most traders quote volatility directly and then calculate the market price of options implied by this volatility quote.

volatility of a risk factor from other related risks, and then construct instruments that can be used to trade it. An important point should be emphasized here. The introductory discussion contained in this chapter deals with the case where the volatility parameter is a function of time and the underlying price only. These methods have to be modified for more complex volatility specifications.

Suggested Reading

Rebonato (2000) and (2002) are good places to start getting acquainted with the various notions of volatility. *Rebonato* (2002) deals with the Libor market model and puts volatility in this context as well. Some of the material in this chapter comes directly from *Demeterfi et al.* (1999), where the reader will find proper references to the literature as well. The important paper by *Dupire* (1992) and the literature it generated can be consulted for local volatility.

Exercises

1. Read the quote carefully and describe how you would take this position using volatility swaps. Be precise about the parameters of these swaps.
 - (a) How would you price this position? What does pricing mean in this context anyway? Which price are we trying to determine and write in the contract?
 - (b) In particular, do you need the correlations between the two markets?
 - (c) Do you need to know the smile before you sell the position?
 - (d) Discuss the risks involved in this volatility position.

Volatility Swaps

A bank is recommending a trade in which investors can take advantage of the wide differential between Nasdaq 100 and S&P500 longer-dated implied volatilities.

Two-year implied volatility on the Nasdaq 100 was last week near all-time highs, at around 45.7%, but the tumult in tech stocks over the last several years is largely played out, said [a] global head of equity derivatives strategy in New York. The tech stock boom appears to be over, as does the most eye-popping part of the downturn, he added. While there will be selling pressure on tech companies over the next several quarters, a dramatic sell-off similar to what the market has seen over the last six months is unlikely.

The bank recommends entering a volatility swap on the differential between the Nasdaq and the S&P, where the investor receives a payout if the realized volatility in two years is less than about 21%, the approximate differential last week between the at-the-money forward two-year implied volatilities on the indices. The investor profits here if, in two years, the realized two-year volatility for the Nasdaq has fallen relative to the equivalent volatility on the S&P.

It might make sense just to sell Nasdaq vol, said [the trader], but it's better to put on a relative value trade with the Nasdaq and S&P to help reduce the volatility beta in the Nasdaq position. In other words, if there is a total market meltdown, tech stocks and the market as a whole will see higher implied vols. But volatility on the S&P500, which represents stocks in a broader array of sectors, is likely to increase substantially, while volatility on the Nasdaq is already close to all-time highs. A relative value trade where the investor takes a view on the differential between the realized volatility in two years time on the two indices allows the investor to profit from a fall in Nasdaq volatility relative to the S&P.

The two-year sector is a good place to look at this differential, said [the trader]. Two years is enough time for the current market turmoil, particularly in the technology sector, to play itself out, and the differential between two-year implied vols, at about 22% last week, is near all-time high levels. Since 1990, the realized volatility differential has tended to be closer to 10.7% over long periods of time.

[The trader] noted that there are other means of putting on this trade, such as selling two-year at-the-money forward straddles on Nasdaq volatility and

buying two-year at-the-money forward straddles on S&P vol. (Derivatives Week, October 30, 2000)

2. The following reading deals with another example of how spread positions on volatility can be taken. Yet, of interest here are further aspects of volatility positions. In fact, the episode is an example of the use of knock-in and knock-out options in volatility positions.
 - (a) Suppose the investor sells short-dated (one-month) volatility and buys six-month volatility. In what sense is this a naked volatility position? What are the risks? Explain using volatility swaps as an underlying instrument.
 - (b) Explain how a one-month break-out clause can hedge this situation.
 - (c) How would the straddles gain value when the additional premium is triggered?
 - (d) What are the risks, if any, of the position with break-out clauses?
 - (e) Is this a pure volatility position?

Sterling volatility is peaking ahead of the introduction of the euro next year. A bank suggests the following strategy to take advantage of the highly inverted volatility curve. Sterling will not join the euro in January and the market expects reduced sterling positions. This view has pushed up one-month sterling/Deutsche mark vols to levels of 12.6% early last week. In contrast, six-month vols are languishing at under 9.2%. This suggests selling short-dated vol and buying six-month vols. Customers can buy a six-month straddle with a one-month break-out clause added to replicate a short volatility position in the one-month maturity. This way they don't have a naked volatility position. (Based on an article in Derivatives Week.)