8

Dynamic Hedging

From our discussion thus far, it ought to be obvious why serious option traders use theoretical pricing models. First, a model tells us something about an option's value. We can compare this value with the price of the option in the marketplace and from this choose an appropriate strategy. Second, once we have taken a position, the model helps us quantify many of the risks that option trading entails. By understanding these risks, we will be better prepared to minimize our losses when market conditions move against us and maximize our profits when market conditions move in our favor.

In discussing the performance of a theoretical pricing model, it is important to remember that all models are probability based. Even if we assume that we have all the right inputs into the model and that the model itself is correct, there is no guarantee that we will show a profit on any one trade. More often than not, the actual results will deviate, sometimes significantly, from what is predicted by the theoretical pricing model. It is only over many trades that the results will even out so that, on average, we achieve a result close to that predicted by the theoretical pricing model.

However, option-pricing theory also suggests that for a single option trade there is a method by which we can reduce the variations in outcome so that the actual results will more closely approximate what is predicted by the theoretical pricing model. By treating the life of an option as a series of bets, rather than one bet, the model can be used to replicate long-term probability theory.

Consider the following situation:

Stock price = $97.70
Time to June expiration = 10 weeks
Interest rate = 6.00 percent

Suppose that we are using a theoretical pricing model to evaluate June options on this stock. We already have three inputs into the model—underlying price, time to expiration, and the interest rate—but we still need three additional inputs—exercise price, type, and volatility. Given that we can choose from among the available exercise prices and that we can also choose the type of option (either call or put), we still lack the one unobservable input—volatility. In theory, we would like to know the future realized volatility of the underlying stock over the next 10 weeks. Clearly, we can never know the future, but let's imagine that we have a crystal ball that can predict the future. When we look into our crystal ball, we see that the volatility of the stock over the next 10 weeks will be 37.62 percent.

The June 100 call, being very close to at the money, is likely to be actively traded, so let's focus on that option. Feeding our inputs into the Black-Scholes model, we find that the June 100 call has a theoretical value of 5.89. When we check its price in the marketplace, we find that it is being offered at 5.00. How can we profit from this discrepancy?

Clearly, our first move will be to purchase the June 100 call because it is underpriced by 0.89. Can we now walk away from the position and come back at expiration to collect our money? In our previous discussion of theoretical pricing models, we noted that the purchase or sale of a theoretically mispriced option requires us to establish a neutral hedge by taking an opposing position in the underlying contract. When this is done correctly, for small changes in the price of the underlying contract, the increase or decrease in the value of the option position will exactly offset the decrease or increase in the value of the opposing position in the underlying contract. Such a hedge is unbiased, or neutral, with respect to directional moves in the underlying contract.

In order to establish the appropriate riskless hedge, we need to determine the delta of the June 100 call. Using our theoretical pricing model, we find that the option has a delta of 50. For each call we purchase, we must sell 0.50, or one-half, of an underlying contract. Because it is usually not possible to buy or sell fractional underlying contracts, let's assume that we buy 100 June 100 calls and sell 50 underlying contracts.[1] We now have the following delta-neutral position:

Position	Contract Delta	Delta Position
Long 100 June 100 calls	50	+5,000
Short 50 underlying contracts	100	−5,000

Suppose that one week later the price of the stock has moved up to 99.50. At this point, we can feed the new market conditions into our theoretical pricing model:

Stock price = 99.50
Interest rate = 6.00 percent
Time to June expiration = 9 weeks
Volatility = 37.62 percent

Note that we have made no change in the interest rate or volatility. Theoretical pricing models typically assume that these two inputs remain constant over the life of the option.[2] Based on the new inputs, we can calculate the new delta for the June 100 call, in this case 54.

[1] The underlying contract for most stock options is 100 shares of stock. The proper hedge is therefore equivalent to selling 5,000 shares of stock.

[2] Whether this is in fact a realistic assumption we will leave for a later discussion.

Position	Contract Delta	Delta Position
Long 100 June 100 calls	54	+5,400
Short 50 underlying contracts	100	–5,000

Our delta position is now +400. We can think of this as the end of one bet, with another bet about to begin.

Whenever we begin a new bet, we are required to return to a delta-neutral position. In our example, it will be necessary to reduce our position by 400 deltas. There are a number of ways to do this, but to keep our present calculations as simple as possible and to remain consistent with the theoretical pricing model, we will make the necessary trades in the underlying contract because an underlying contract always has a delta of 100. We can return to delta neutral by selling 4 underlying contracts. Our position is now

Position	Contract Delta	Delta Position
Long 100 June 100 calls	54	+5,400
Short 54 underlying contracts	100	–5,400

We are again delta neutral and about to begin a new bet. As before, our new bet depends only on the volatility of the underlying contract, not its direction.

The extra four underlying contracts that we sold were an *adjustment* to our position. In option trading, adjustments are trades that are made primarily to ensure that a position remains delta neutral. In our case, the sale of the four extra contracts has no effect on our theoretical edge because, from an option trader's point of view, an underlying contract has no theoretical value. The trade is made solely for the purpose of adjusting our hedge to remain delta neutral.

In Chapter 17, we will look at the use of options to protect a preexisting position. Such protective strategies usually employ a *static hedge*, whereby opposing market positions are taken in different contracts, with the entire position being carried to a fixed maturity date. To capture an option's mispricing, the theoretical pricing model requires us to employ a *dynamic hedging* strategy. We must periodically reevaluate the position to determine the delta of the position and then buy or sell an appropriate number of underlying contracts to return to delta neutral. This procedure must be followed over the entire life of the option.

Because volatility is assumed to compound continuously, theoretical pricing models assume that adjustments are also made continuously and that the hedge is being adjusted at every moment in time. Such continuous adjustments are not possible in the real world because a trader can only trade at discrete intervals. By making adjustments at regular intervals, we are conforming as closely as possible to the principles of the theoretical pricing model.

The entire dynamic hedging process for our hedge, with adjustments made at weekly intervals, is shown in Figure 8-1. At the end of each interval, the delta of the June 100 call was recalculated from the time remaining to expiration, the current price of the underlying contract, an interest rate of 6.00 percent, and a volatility of 37.65 percent. Note that we did not change the volatility,

Figure 8-1

Stock price = 97.70			Time to June expiration = 10 weeks Interest rate = 6.00% Volatility = 37.62%				
June 100 call:			Price = 5.00 (Implied volatility = 32.40%) Theoretical value = 5.89 Delta = 50				
Week	Share Price	Delta of 100 Call	Total Delta Position	Adjustment (Contracts)	Total Adjustments	Adjustment Cash Flow	Interest on Adjustments
0	97.70	50	0				
1	99.50	54	+400	Sell 4	Short 4	+398.00	+4.12
2	92.75	35	−1900	Buy 19	Long 15	−1762.25	−16.22
3	95.85	43	+800	Sell 8	Long 7	+766.80	+6.18
4	96.20	43	0	None	Long 7	0	0
5	102.45	62	+1900	Sell 19	Short 12	+1946.55	+11.20
6	93.30	28	−3400	Buy 34	Long 22	−3172.20	−14.60
7	91.15	17	−1100	Buy 11	Long 33	−1002.65	−3.46
8	95.20	27	+1000	Sell 10	Long 23	+952.00	+2.19
9	102.80	72	+4500	Sell 45	Short 22	+4626.00	+5.32
10	103.85			Buy 22		−2284.70	

even though other market conditions may have changed. Volatility, like interest rates, is assumed to be constant over the life of the option.[3]

What will we do with our position at the end of 10 weeks when the options expire? At that time, we plan to close out the position by

1. Letting any out-of-the-money options expire worthless
2. Selling any in-the-money options at parity (intrinsic value) or, equivalently, exercising them and offsetting them against the underlying contract
3. Liquidating any outstanding underlying contracts at the market price

Let's go through this procedure step by step and see what the complete results of our hedge are.

Original Hedge

At June expiration (week 10), with the underlying contract at 103.85, we can close out the June 100 calls by either selling them at 3.85 or exercising the calls

[3] In practice, as new information becomes available, traders are constantly changing their opinions about interest rates and volatility. Here we make the assumption of constant volatility and interest rates in order to be consistent with option pricing theory.

and selling the underlying contract. Either method will result in a credit of 3.85 to our account. Because we originally paid 5.00 for each option, we will show a loss on our option position of

$$100 \times (3.85 - 5.00) = 100 \times -1.15 = -115.00$$

As part of our original hedge, we also sold 50 underlying contracts at 97.70. At expiration, in order to close out the position, we were required to buy them back at 103.85, for a loss of 6.15 per contract. Our total loss on the underlying trade is therefore

$$50 \times (97.70 - 103.85) = 50 \times -6.15 = -307.50$$

Adding this to our option loss, the total loss on the original hedge is

$$-115.00 - 307.50 = -422.50$$

This certainly does not appear to have been successful. We expected to make money on the position, yet it appears that we have a sizable loss.

Adjustments

Fortunately, the original hedge was not our only transaction. In order to remain delta neutral over the 10-week life of the option, we were forced to buy and sell underlying contracts. At the end of week 1, we were long 400 deltas, so we were required to sell four underlying contracts at 99.50. At the end of week 2, we were short 1,900 deltas, so we were required to buy 19 underlying contracts at 92.75, and so on each week until the end of week 10. At expiration, with the underlying contract at 103.85, we bought in the 22 underlying contracts that we were short at the end of week 9.

In this example, each time the underlying price rose, our delta position became positive, so we were forced to sell underlying contracts, and each time the underlying price fell, our delta position became negative, so we were forced to buy underlying contracts. Because our adjustments depended only on our delta position, we were forced to do what every trader wants to do: buy low and sell high.

The result of making all the adjustments required to maintain a delta-neutral position was a profit of 467.55. (The reader may wish to confirm this by adding up the cash flow from all the trades in the adjustment column in Figure 8-1.) This profit more than offset the losses incurred from the original hedge.

Interest Lost on the Option Position

We originally bought 100 June options at a price of 5.00 each, for a total cash outlay of 500.00. At the assumed interest rate of 6.00 percent, the cost of financing the option purchase for the 10-week (70-day) life of the position was

$$-500.00 \times 6\% \times 70/365 = -5.75$$

Interest Earned on the Stock Position

To establish our initial hedge, we sold 50 underlying stock contracts at a price of 97.70 each, for a total credit of 4,885.00. Over the life of the hedge, we were able to earn total interest in the amount of

$$+4{,}885 \times 6\% \times 70/365 = +56.21$$

Interest on the Adjustments

Each week we were forced to buy or sell underlying contracts in order to remain delta neutral. As a result, there was either a cash debit on which we were required to pay interest or a cash credit on which we were able to earn interest. For example, at the end of week 1, we were forced to sell four underlying contracts at a price of 99.50 each, for a total credit of 4 × 99.50 = 398.00. The interest earned on this credit for the remaining nine weeks was

$$+398.00 \times 6\% \times 63/365 = +4.12$$

At the end of week 2, we were forced to buy 19 underlying contracts at a price of 92.75 each, for a total debit of 19 × 92.75 = 1,762.25. The interest cost on this debit over the remaining eight weeks was

$$-1{,}762.25 \times 6\% \times 56/365 = -16.22$$

Adding up the interest on all the adjustments, we get a total of –5.28.

Dividends

To keep our example relatively simple, we have assumed that the stock pays no dividend over the life of the option. If the stock were to pay a dividend, any long stock position resulting from either the original hedge or the adjustment process would receive the dividend. Any short stock position would be required to pay out the dividend. There also would be an interest consideration on the amount of the dividend, interest either earned or interest lost, between the date of the dividend payment and expiration. The dividend and the interest on the dividend would then become part of the total profit or loss.

What was the total cash flow resulting from the entire 10-week hedge? This amount, +90.24, is shown in Figure 8-2. Of course, this represents the cash flow at the end of 10 weeks. To obtain the initial or present value, we need to discount backwards over 10 weeks at an interest rate of 6.00 percent. This gives us a final value, or total profit and loss (P&L), of

$$\frac{90.24}{1+0.06\times 70/365} = 89.21$$

How does this final value of 89.21 compare with our predicted profit or loss? We purchased 100 June options at a price of 5.00 each, but the options had a theoretical value of 5.89, so the theoretical profit was

$$100 \times (5.89 - 5.00) = +89.00$$

Figure 8-2

Dynamic Hedging Results		
Original hedge P&L:		−422.50
Option P&L	100 × (3.85 − 5.00) = −115.00	
Stock P&L	50 × (97.70 − 103.85) = −307.50	
Adjustment P&L:		+467.55
Carry (interest) on the options:		
	100 × −5.00 × 6.00% × 70/365 = −5.75	−5.75
Carry (interest) on the stock:		
	50 × +97.70 × 6.00% × 70/365 = +56.21	+56.21
Interest on the adjustments:		−5.27
Total cash flow:		+90.24
Discounted cash flow:	90.24/(1 + 0.06 × 70/365) = 89.21	+89.21
Predicted P&L:	100 × (5.89 − 5.00) = 100 × 0.89 = 89.00	+89.00

In our example, the profit and loss were made up of five components. Two of these were positive (the adjustments and the interest earned on stock), while three were negative (the original hedge, the option carrying costs, and interest on the adjustments). Is this always the case? Because price movement in the underlying contract is assumed to be random, it is impossible to determine beforehand which components will be profitable and which will not. It would also be possible to construct an example where the original hedge was profitable and the adjustments were not. The important point is that if a trader's inputs are correct, in some combination, he can expect to show a profit or loss approximately equal to that predicted by the theoretical pricing model.

Of all the inputs, volatility is the only one that is not directly observable. Where did our volatility figure of 37.62 percent come from? Obviously, it is not possible to know the future volatility. In our example the 10 price changes in Figure 8-1 do in fact represent an annualized volatility of 37.62 percent. The complete calculations are given in Appendix B.

In the foregoing example, we assumed that the market was *frictionless*, that no external factors affected the total profit or loss. This assumption is basic to many financial models. In a frictionless market, we assume that

1. Traders can freely buy or sell the underlying contract without restriction.
2. Traders can borrow and lend as much money as desired at one constant interest rate.

3. Transaction costs are zero.
4. There are no tax consequences.

A trader will immediately realize that option markets are not frictionless because in the real world, each of these assumptions is violated to a greater or lesser degree. In our example, we were required to sell stock to initiate the original hedge. If we did not own the stock, we would need to *sell short* by first borrowing the stock and then making delivery. In some markets, short sales may be difficult to execute because of exchange or regulatory restrictions. Moreover, even if a short sale is possible, a trader typically will not receive full interest on the proceeds from the short sale.

Turning to options on futures, in some markets, there is a daily limit on the amount of allowable price movement for a futures contract. When this limit is reached, the market is *locked*, and no further trading can take place until the price of the futures contract comes off its limit. Clearly, in such markets, the underlying contract cannot always be freely bought or sold.

Concerning interest rates, different rates apply to different market participants. The rate that applies to an individual trader will not be the same rate that applies to a large financial institution. Moreover, even for the same trader, different rates can apply to different transactions. If a trader has a debit balance, it will cost him more to carry that debit; if he has a credit balance, he will not earn as much on that credit. There is a spread, and perhaps a fairly large one, between a trader's borrowing and lending rate. Fortunately, the interest-rate component is usually the least important of the inputs into a theoretical pricing model. Even though the applicable interest rate may vary from trader to trader, in general, it will cause only minor changes in the total profit or loss in relation to the profit or loss resulting from other inputs.

Transaction costs, on the other hand, can be a very real consideration. If these costs are high, the hedge in Figure 8-1 might not be a viable strategy; all the profits could be eaten up by brokerage and exchange fees. The desirability of a strategy will depend not only on the trader's initial transaction costs but also on the subsequent costs of making adjustments. The adjustment cost is a function of a trader's desire to remain delta neutral. A trader who wants to remain delta neutral at every moment will have to adjust more often, and more adjustments mean greater transaction costs.

If a trader initiates a hedge but adjusts less frequently or does not adjust at all, how will this affect the outcome? Because theoretical evaluation of options is based on the laws of probability, a trader who initiates a theoretically profitable hedge still has the odds on his side. Although he may lose on any one individual hedge, if given a chance to initiate the same hedge repeatedly at a positive theoretical edge, on average, he should profit by the amount predicted by the theoretical pricing model. The adjustment process is simply a way of smoothing out the winning and losing hedges by forcing the trader to make more bets, always at the same favorable odds. A trader who is disinclined to adjust is at greater risk of not realizing a profit on any one hedge. Adjustments do not in themselves alter the expected return; they simply reduce the short-term effects of good and bad luck.

Based on the foregoing discussion, a retail customer and a professional trader are likely to approach option trading in a somewhat different manner, even if both understand and use the values generated by a theoretical pricing model. A professional trader, particularly if he is an exchange member, has relatively low transaction costs. Because adjustments cost him very little in relation to the expected theoretical profit from a hedge, he will be inclined to make frequent adjustments. In contrast, a retail customer who establishes the same hedge will be less inclined to adjust or will adjust less frequently because any adjustments will reduce the profitability of the position. A retail customer who understands the laws of probability will realize that his position has the same favorable odds as the professional trader's position, but he should also realize that his position is more sensitive to the effects of short-term good and bad luck. Even though the retail customer may occasionally experience larger losses than the professional trader, he will also occasionally experience larger profits. In the long run, on average, both should end up with approximately the same profit.[4]

Taxes may also be a factor in evaluating an option strategy. When positions are initiated, when they are liquidated, how the positions overlap, and the relationship between different instruments (e.g., options, stock, futures, physical commodities, etc.) may have different tax consequences. Such consequences may affect the value of a diversified portfolio, and for this reason, portfolio managers must be sensitive to the tax ramifications of a strategy. Because each trader has unique tax considerations and this book is intended as a general guide to option evaluation and strategies, we will simply assume that each trader wishes to maximize his pretax profits and that he will worry about taxes afterward.

It may seem like a fortunate coincidence that the theoretical P&L in our example and the actual P&L are so close. In fact, the example in Figure 8-1 was carefully constructed to demonstrate why the dynamic hedging process is so important. In the real world, it is unlikely that the actual results from any one hedge will so closely match the theoretical results.

Figure 8-3 illustrates in graphic terms the dynamic hedging process. We determined the initial delta of the option (the dotted line) at the underlying price of 97.70 and then took an opposing delta position in the underlying contract (the dashed line). For very small moves in the underlying price, the profit from one position offset the loss from the other position. As the changes in the underlying price in either direction become greater, because of the option's curvature (its gamma), there is a mismatch between these two positions. With a falling underlying price, the rate at which the option position loses value begins to decline; with a rising underlying price, the rate at which the option position gains value begins to increase. In Figure 8-3*a*, we can see this mismatch, or unhedged amount, at an underlying price of 99.50.

With the underlying price at 99.50, we captured the value of this mismatch by adjusting the position to return to delta neutral. This is shown in Figure 8-3*b*. We recalculated the delta at the new underlying price and took a new opposing position in the underlying contract. When the underlying price fell to 92.75, there was again a mismatch equal to the unhedged amount.

[4] This, of course, ignores the very real advantage the professional trader often has from being able to buy at the bid price and sell at the ask price. A retail customer can never hope to match the profit resulting from this advantage, nor should he try to do so.

Figure 8-3

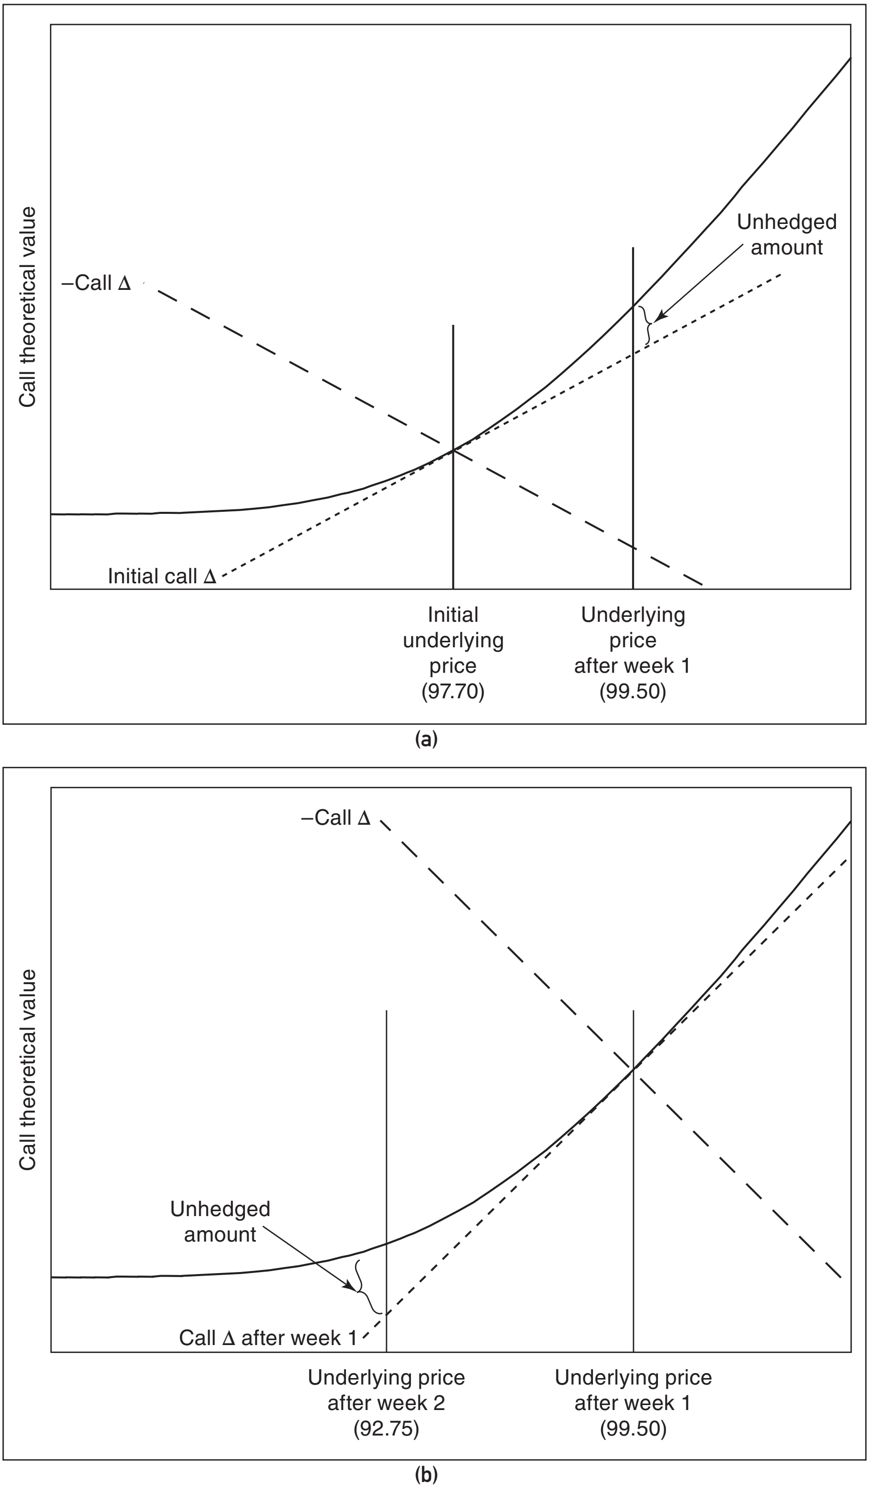
Call theoretical value
–Call Δ
Unhedged amount
Initial call Δ
Initial underlying price (97.70)
Underlying price after week 1 (99.50)
(a)
Call theoretical value
–Call Δ
Unhedged amount
Call Δ after week 1
Underlying price after week 2 (92.75)
Underlying price after week 1 (99.50)
(b)

By rehedging the position each week, we were able to capture a series of profits resulting from the mismatch between the option's changing delta and the fixed delta of the underlying contract. Of course, while time was passing, there were also interest considerations. But most of the option's value was determined by the amount earned on the rehedging process. In theory, if we ignore interest, the sum of all these small profits (the unhedged amounts in Figure 8-3) should approximately equal the value of the option

$$\text{Option theoretical value} = \{\cdot\} + \{\cdot\} + \{\cdot\} + \cdots + \{\cdot\} + \{\cdot\} + \{\cdot\}$$

In our example, rehedging took place at discrete intervals, equivalent to making a finite number of bets, all with the same positive theoretical edge. If we want to exactly replicate the option's theoretical value, we need to make an infinite number of bets. This, in theory, can only be achieved by continuous rehedging of the position at every possible moment in time. If such a process were possible, and if all the assumptions on which the model is based were accurate, then the rehedging process would indeed replicate the exact value of the option.

Of course, continuous rehedging is not possible in the real world. Nor are all the model assumptions entirely accurate. Nonetheless, most traders have found through experience that using a dynamic hedging strategy, even if only at discrete intervals, is the best way to capture the difference between an option's price and its theoretical value.

Given that continuous rehedging is not possible, how often should a trader rehedge? The answer to this question will depend on each trader's cost structure and risk tolerance. We have already noted that a trader's transaction costs are likely to affect the frequency at which adjustments are made. Higher transaction costs will often lead to less frequent adjustments. If we ignore the question of transaction costs, there are two common approaches to rehedging: rehedge at regular intervals or rehedge whenever the delta becomes unbalanced by a predetermined amount.

The position in Figure 8-1 is an example of the first approach—rehedging at regular intervals. Here we made adjustments to the position at the end of each week. Of course, we might have made adjustments at the end of each day or even every hour if we were willing to recalculate the deltas so frequently. The more often rehedging takes place, the more likely it is that the final result will approximate the results predicted by the model. In our example, we used weekly intervals for no other reason than 10 lines seemed to fit nicely on the page.

Most traders do not insist on maintaining an exactly delta-neutral position. Within limits, they are willing to accept some directional risk. The more directional risk a trader is willing to accept, the less frequent the adjustments. And the less frequent the adjustments, the more likely the actual results will differ from the results predicted by the theoretical pricing model. For example, if a trader decides that he is willing to accept a directional risk up to 500 deltas, no rehedging would take place after week 1 (+400 deltas). If the trader is willing to accept a directional risk up to 1,000 deltas, no rehedging would take place at the end of week 1 (+400 deltas), week 3 (+800 deltas), and week 8 (+1,000 deltas). And if the trader is willing to accept a directional risk up to 1,500 deltas,[5] no rehedging would take place at the end of week 1 (+400 deltas), week 3 (+800 deltas), week 7 (–1,100 deltas), and week 8 (+1,000 deltas). In each case, because of the less frequent rehedging, the actual results are more likely to differ from the predicted results.

Note that after the hedge in Figure 8-1 was initiated, no subsequent trades were made in the option market. The trader's only concern was the realized volatility, or price fluctuations, in the underlying market. These price fluctuations determined the size and frequency of the adjustments, and in the final analysis, it was the adjustments that determined the profitability of the hedge. We might think of the hedge as a race between the loss in time value of the June 100 calls and the cash flow resulting from the adjustments, with the theoretical pricing model acting as the judge. Under the assumptions of the model, if options are purchased at less than theoretical value, the adjustments will win the race; if options are purchased at more than theoretical value, the loss in time value will win the race. The conditions of the race are determined by the inputs into the theoretical pricing model.

We assumed in our example that the future volatility was known to be 37.62 percent. What will be the outcome if volatility is something other than 37.62 percent? Suppose, for example, that volatility turns out to be higher than 37.62 percent. Higher volatility means greater price fluctuations, resulting in more and larger adjustments. In our example, more adjustments mean more profit. This is consistent with the principle that higher volatility increases the value of options.

What about the opposite, if volatility is less than 37.62 percent? Lower volatility means smaller price fluctuations, resulting in fewer and smaller adjustments. This will reduce the profit. If the volatility is low enough, the adjustment profit will just offset the other components, so the total profit from the hedge will be exactly zero. This *breakeven volatility* is identical to the option's implied volatility at the original trade price. Using the Black-Scholes model, we find that the implied volatility of the June 100 call at a price of 5.00 is 32.40 percent. At this volatility, the race between profits from the adjustments and the loss in the option's time value will end in an exact tie. Above a volatility of 32.40 percent, we expect the hedge, including adjustments and interest, to show a profit; below 32.40 percent, we expect the hedge to show a loss.

Because we needed to make adjustments to realize a profit, it may seem that every profitable hedge requires us to maintain the position until expiration. In practice, this may not be necessary. Suppose that immediately after we establish the hedge, the implied volatility in the option market increases from 32.40 percent, the implied volatility at which we bought the June 100 calls, to 37.62 percent, the realized volatility of the underlying contract we expect over the life of the option. What will happen to the price of the June 100 call? Its price will rise from 5.00 (an implied volatility of 32.40 percent) to 5.89 (an implied volatility of 37.62 percent). We can then sell our calls for an immediate profit of 0.89 per option. Of course, if we want to close out the hedge,

[5] These delta numbers were chosen only to illustrate the effect of rehedging based on a predetermined delta risk. Even a directional risk of 500 deltas might be more than many traders are willing to accept.

we must also buy back the 50 underlying contracts that we originally sold. What effect will the change in implied volatility have on the price of these contracts? Implied volatility is a characteristic associated with options, not with underlying contracts. Consequently, we expect the underlying contract to continue to trade at its original price of 97.70. By purchasing our 50 outstanding underlying contracts at a price of 97.70, we will realize an immediate total profit from the hedge of 89.00, exactly the amount predicted by the theoretical pricing model. If we can do all this, there is no reason to hold the position for the full 10 weeks.

How likely is an immediate reevaluation of implied volatility from 32.40 to 37.62 percent? Although swift changes in implied volatility occur occasionally, more often changes occur gradually over a period of time and are the result of equally gradual changes in the volatility of the underlying contract. As the volatility of the underlying contract changes, option demand rises and falls, and this demand is reflected in a corresponding rise or fall in the implied volatility. In our example, if the price of the underlying contract begins to fluctuate at a volatility greater than 32.40 percent, we can expect implied volatility to rise. If implied volatility ever reaches our target of 37.62 percent, we can simply sell our calls and buy our underlying contracts, thereby realizing our expected profit of 89.00 without having to hold the position for the full 10 weeks. But option prices are subject to a wide variety of market forces, not all of them theoretical. There is no guarantee that implied volatility will ever reevaluate upward to 37.62 percent. In this case, we will have to hold the position and continue to adjust for the full 10 weeks to realize our profit.

Every trader hopes that implied volatility will reevaluate as quickly as possible toward his volatility target. It not only enables him to realize his profits more quickly, but it eliminates the risk of holding a position for an extended period of time. The longer a position is held, the greater the possibility of error from the inputs into the model.

Not only might implied volatility not reevaluate favorably, it also might move against us, even if the actual volatility of the underlying contract moves in our favor. Suppose that after initiating our hedge, implied volatility immediately falls from 32.40 to 30.35 percent. The price of the June 100 call will fall from 5.00 to 4.65, and we will have an immediate loss of $100 \times -0.35 = -35.00$. Does this mean that we made a bad trade and should close out the position? If the volatility forecast of 37.65 percent turns out to be correct, the options will still be worth 5.89 by expiration. If we hold the position and adjust, we can eventually expect a profit of 89.00 points. Realizing this, we ought to maintain the position as we originally intended. Even though an adverse move in implied volatility is unpleasant, it is something with which all traders must learn to cope. Just as a speculator can rarely hope to pick the exact bottom or top at which to take a long or short position, an option trader can rarely hope to pick the exact bottom or top in implied volatility. He must try to establish positions when market conditions are favorable. But he must also realize that conditions might become even more favorable. If they do, his initial trade may show a temporary loss. This is something a trader learns to accept as a practical aspect of trading.

Let's look at one other dynamic hedging example, this time in the form of an overpriced put in the futures option market. Suppose that current market conditions are as follows:

Futures price = 61.85
Time to March expiration = 10 weeks
Interest rate = 8.00 percent

Again, let's assume that we know the true volatility of the underlying contract over the 10-week life of the option, in this case 21.48 percent. In this example, we will focus on the March 60 put, with a theoretical value of 1.46 but a price of 1.70, equivalent to an implied volatility of 23.92 percent.

Because the put is overpriced, we will begin by selling 100 March 60 puts, with a delta of –35 each, and simultaneously selling 35 underlying futures contracts. We will then follow a dynamic hedging procedure by recalculating the put delta at the end of each week and buying or selling futures to remain delta neutral. At expiration, we will close out the entire position. The entire dynamic hedging process is shown in Figure 8-4.

The cash flow in this example is slightly different from that in our stock option example. Although these are options on futures contracts and in many markets are subject to futures-type settlement, we will follow the U.S. convention and assume that the options are subject to stock-type settlement, requiring immediate and full cash payment. Futures, however, are always subject to futures-type settlement: there is no initial cash outlay, but a cash flow, in the form of variation, will result whenever the price of the futures contract changes. When this occurs, there will be a cash credit, on which interest can be earned, or a cash debit, on which interest must be paid.

Figure 8-4

Futures price = 61.85			Time to March expiration = 10 weeks Interest rate = 8.00% Volatility = 21.48%				
March 60 put:			Price = 1.70 (Implied volatility = 23.92%) Theoretical value = 1.46 Delta = -35				
Week	Share Price	Delta of 60 Put	Total Delta Position	Adjustment (Contracts)	Total Adjustments	Variation	Interest on Variation
0	61.85	–35	0				
1	60.83	–42	+700	Sell 7	Short 7	+35.70	+0.49
2	62.78	–28	–1400	Buy 14	Long 7	–81.90	–1.01
3	63.16	–24	–400	Buy 4	Long 11	–10.64	–0.11
4	61.68	–34	+1000	Sell 10	Long 1	+35.52	+0.33
5	59.86	–50	+1600	Sell 16	Short 15	+61.88	+0.47
6	62.88	–21	–2900	Buy 29	Long 14	–151.00	–0.93
7	61.50	–31	+1000	Sell 10	Long 4	+29.98	+0.13
8	62.60	–15	–1600	Buy 16	Long 20	–34.10	–0.10
9	60.18	–45	–3000	Sell 30	Short 10	+36.30	+0.06
10	58.61			Buy 10			

All P&L components for this example are shown in Figure 8-5. Three of these components are the same as in the stock option example: the P&L on the original hedge, the P&L resulting from the delta-neutral dynamic hedging process, and the carrying cost on the options. However, the interest on the initial stock position, as well as the interest on the adjustments, has been replaced by the interest on the variation credits and debits.

For example, as part of our original hedge, we sold 35 futures contracts at a price of 61.85. After week 1, the futures price declined to 60.83. As a result, we received a variation payment of

$$35 \times (61.85 - 60.83) = 35.70$$

We were able to earn 8.00 percent on this amount for the nine weeks (63 days) remaining to expiration

$$35.70 \times 8\% \times 63/365 = 0.49$$

At the end of week 1, in order to remain delta neutral, we were forced to sell seven futures contracts. This, together with our initial sale of 35 futures, left us short a total of 42 futures. After week 2, the futures price rose to 62.78. The result was a variation debit of

$$42 \times (60.83 - 62.78) = -81.90$$

In order to finance this debit for the eight weeks (56 days) remaining to expiration, we incurred an interest cost of

$$-81.90 \times 8\% \times 56/365 = -1.01$$

The total interest on all variation cash flows was –0.67.

Figure 8-5

Dynamic Hedging Results		
Original hedge P&L:		+144.40
Option P&L	100 × (1.70 – 1.39) = +31.00	
Futures P&L	35 × (61.85 – 58.61) = +113.40	
Adjustment P&L:		–122.01
Carry (interest) on the options:		
	100 × +1.70 × 8.00% × 70/365 = +2.61	+2.61
Interest on the variation:		–0.67
Total P&L:		+24.33
Discounted cash flow:	24.33/(1 + 0.08 × 70/365) = 23.96	+23.96
Predicted P&L:	100 × (1.70 – 1.46) = 100 × 0.24 = 24.00	+24.00

The total cash flow of 24.33 and the present value of this amount, 23.96, are shown in Figure 8-5. The predicted theoretical profit was

$$100 \times (1.70 - 1.46) = 24.00$$

In both our stock option and futures option examples, we were able use the dynamic hedging process to capture the difference between the option's theoretical value and its price. In a sense, dynamic hedging enabled us to take the other side of the trade, but at the option's true theoretical value. When we bought calls in our stock option example, we sold the same calls at theoretical value through the dynamic hedging process. When we sold puts in our futures option example, we bought the same puts at theoretical value through the dynamic hedging process. From this, we can deduce an important principle of option evaluation:

> *In theory, we can replicate an option position through a dynamic hedging process. The cost of this replication is equal to the sum of all the cash flows resulting from the dynamic hedging process. The present value of this sum is equal to the option's theoretical value.*