6

Volatility

What is volatility, and why is it so important in option evaluation? The option trader, like a trader in the underlying instrument, is interested in the direction of the market. But unlike a trader in the underlying, an option trader is also sensitive to the speed of the market. If the market for an underlying contract fails to move at a sufficient speed, options on that contract will have less value because of the reduced likelihood of the market going through an option's exercise price. In a sense, volatility is a measure of the speed of the market. Markets that move slowly are low-volatility markets; markets that move quickly are high-volatility markets.

One might guess intuitively that some markets are more volatile than others. During 2008, the price of crude oil began the year at $99 per barrel, reached a high of $144 per barrel in July, and finished the year at $45 per barrel. The price rose 58 percent and then dropped 69 percent. Yet few traders could imagine a major stock index such as the Standard and Poor's (S&P) 500 Index exhibiting similar fluctuations over a single year.

If we know whether a market will be relatively volatile or relatively quiet and can convey this information to a theoretical pricing model, any evaluation of options on that market will be more accurate than if we simply ignore volatility. Because option models are based on mathematical formulas, we will need some method of quantifying this volatility component so that we can feed it into the model in numerical form.

Random Walks and Normal Distributions

Consider for a moment the pinball maze pictured in Figure 6-1. When a ball is dropped into the maze at the top, it falls downward, pulled by gravity through a series of nails. When the ball encounters each nail, there is a 50 percent chance that the ball will move to the left and a 50 percent chance that it will move to the right. The ball then falls down a level where it encounters another nail. Finally, at the bottom of the maze, the ball falls into one of the troughs.

Figure 6-1 Random walk.

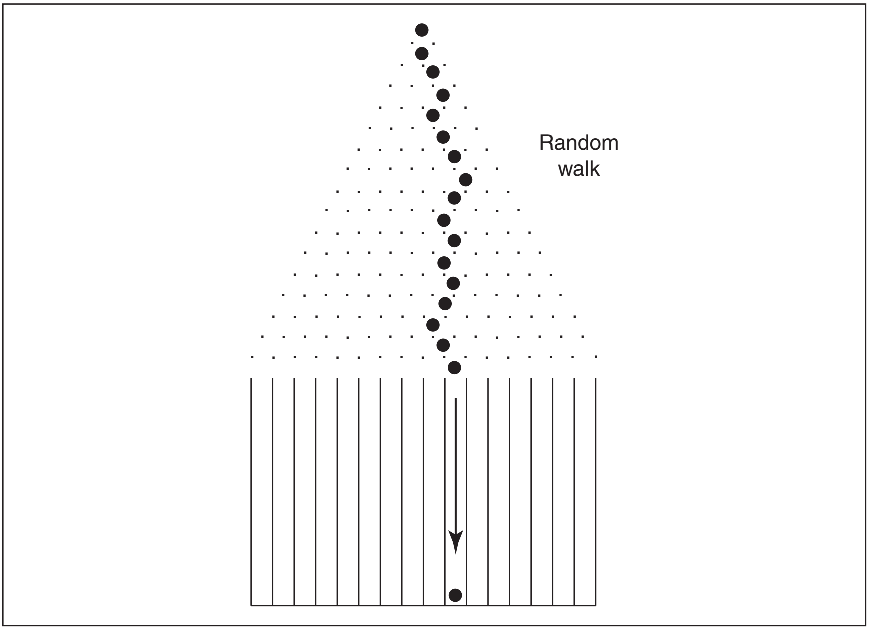

As the ball falls down through the maze, it follows a *random walk*. Once the ball enters the maze, nothing can be done to artificially alter its course, nor can one predict the path that the ball will follow through the maze.

As more balls are dropped into the maze, they might begin to form a distribution similar to that in Figure 6-2. Most of the balls tend to cluster near the center of the maze, with a decreasing number of balls ending up in troughs farther away from the center. If many balls are dropped into the maze, they will begin to form a bell-shaped or *normal distribution*.

If an infinite number of balls were dropped into the maze, the resulting distribution might be approximated by a *normal distribution curve* such as the one overlaid on the distribution in Figure 6-2. Such a curve is symmetrical (if we flip it from right to left, it looks the same), it has its peak in the center, and its tails always move down and away from the center.

Normal distribution curves are used to describe the likely outcomes of random events. For example, the curve in Figure 6-2 might also represent the results of flipping a coin 15 times. Each outcome, or trough, represents the number of heads that occur from each 15 flips. An outcome in trough 0 represents 0 heads and 15 tails; an outcome in trough 15 represents 15 heads and 0 tails. Of course, we would be surprised to flip a coin 15 times and get all heads or all tails. Assuming that the coin is perfectly balanced, some outcome in between, perhaps 8 heads and 7 tails, or 9 heads and 6 tails, seems more likely.

Suppose that we rearrange the nails in our maze so that each time a ball encounters a nail and moves either left or right, it must drop down two levels before it encounters another nail. If we drop enough balls into the maze, we may end up with a distribution similar to the curve in Figure 6-3. Because the

Figure 6-2 Normal distribution.

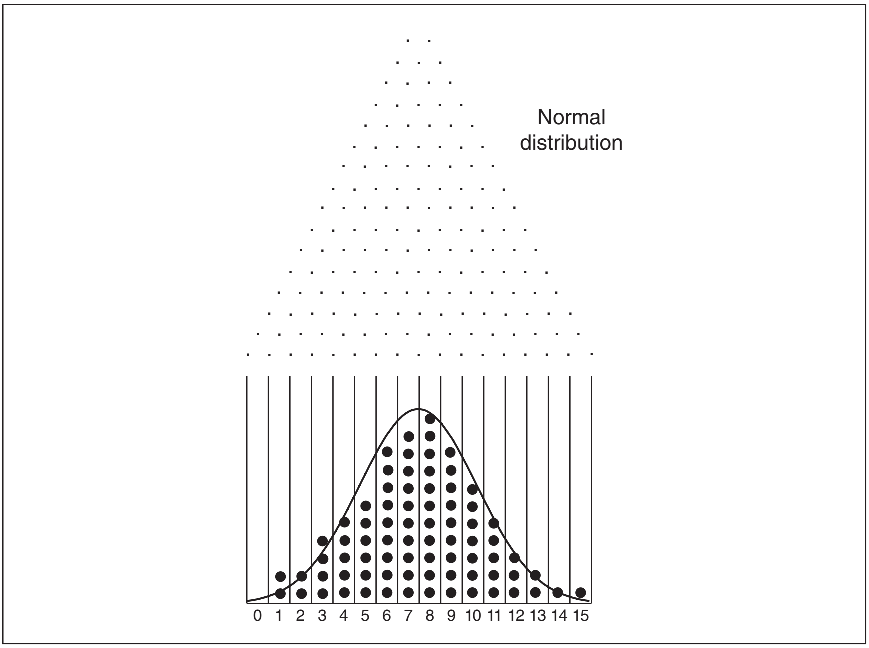

Figure 6-3

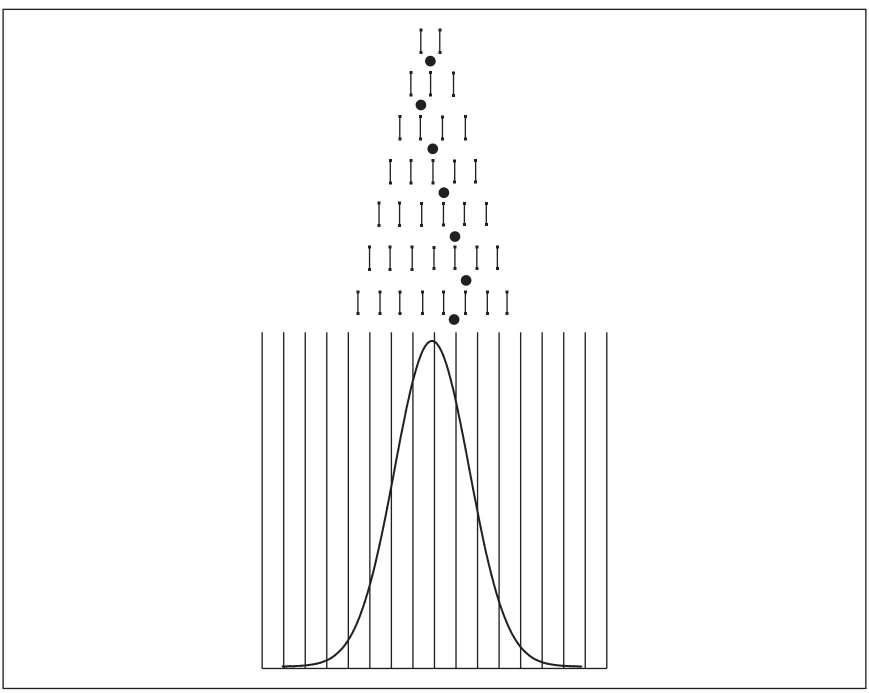

sideways movement of the balls is restricted, the curve will have a higher peak and narrower tails than the curve in Figure 6-2. In spite of its altered shape, the distribution is still normal, although one with slightly different characteristics.

Finally, we might again rearrange the nails so that each time a ball drops down a level, it must move two nails left or right before it can drop down to a new level. If we drop enough balls into the maze, we may get a distribution that resembles the curve in Figure 6-4. This distribution, although still normal, will have a much lower peak and spread out much more quickly than the distributions in either Figure 6-2 or Figure 6-3.[1]

Suppose that we now think of the ball's sideways movement as the up and down price movement of an underlying contract and the ball's downward movement as the passage of time. If the price movement of an underlying contract follows a random walk, the curves in Figures 6-2 through 6-4 might represent possible price distributions in a moderate-, low-, and high-volatility market, respectively.

Earlier in this chapter we suggested that the theoretical pricing of options begins by assigning probabilities to the various underlying prices. How should these probabilities be assigned? One possibility is to assume that, at expiration, the underlying prices are normally distributed. Given that there are many different normal distributions, how will our choice of distribution affect option evaluation?

Because all normal distributions are symmetrical, it may seem that the choice of distribution is irrelevant. Increased volatility may increase the likelihood of large upward movement, but this should be offset by the greater likelihood of large downward movement. However, there is an important distinction between

Figure 6-4

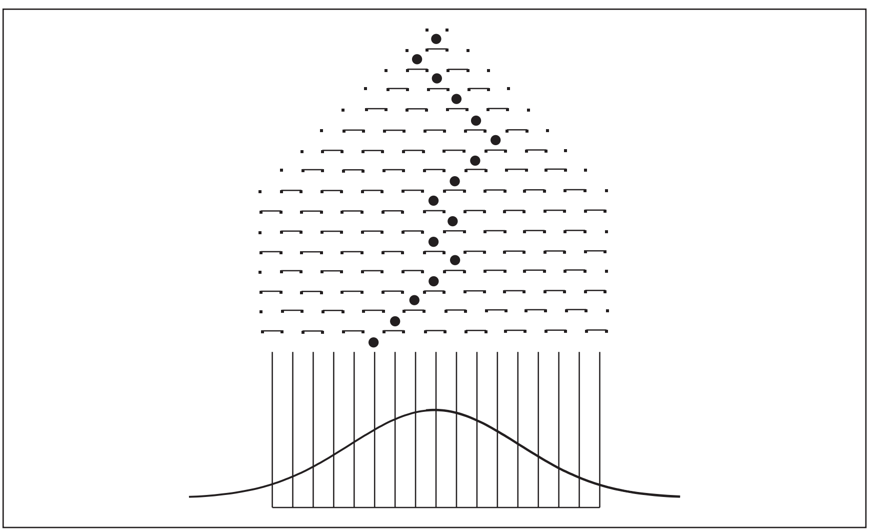

[1] The pinball maze, or *quincunx* (sometimes also called a *Galton board*), pictured in these examples is often used to demonstrate basic probability theory. Examples of a quincunx in action can be found at the following websites:

http://www.teacherlink.org/content/math/interactive/flash/quincunx/quincunx.html
http://www.mathsisfun.com/data/quincunx.html
http://www.jcu.edu/math/isep/Quincunx/Quincunx.html

Figure 6-5

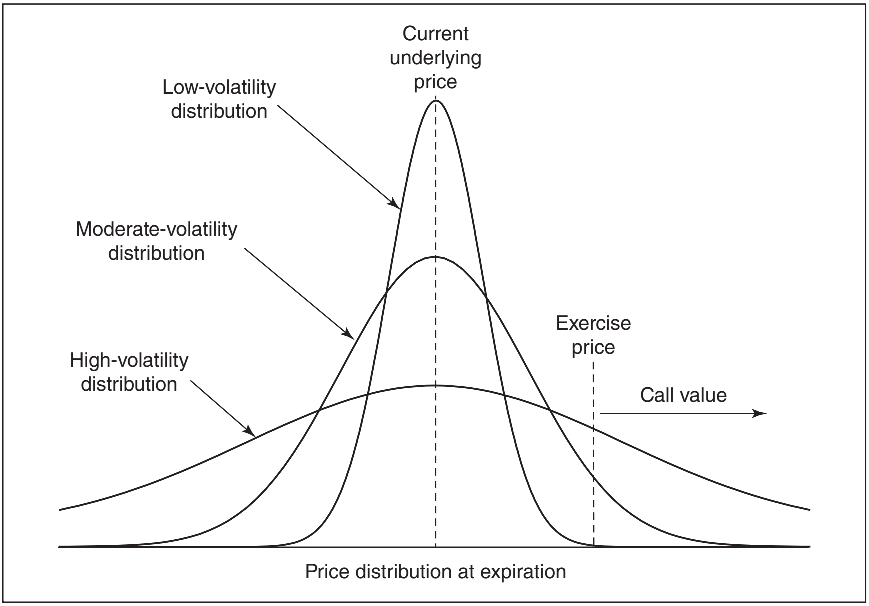

an option position and an underlying position. The expected value for an underlying contract depends on all possible price outcomes. The expected value for an option depends only on the outcomes that result in the option finishing in the money. Everything else is zero.

In Figure 6-5, we have three possible price distributions centered around the current price of an underlying contract. Suppose that we want to evaluate a call at a higher exercise price. The value of the call will depend on the amount of the distribution to the right of the exercise price. We can see that as we move from a low-volatility distribution, to a moderate-volatility distribution, to a high-volatility distribution, a greater portion of the possible price distribution lies to the right of the exercise price. Consequently, the option takes on an increasingly greater value.

We might also consider the value of a put at a lower exercise price. If we assume that movement is random, the same high-volatility distribution that will cause the call to take on greater value will also cause the put to take on greater value. In the case of the put, more of the distribution will lie to the left of the exercise price. Because our distributions are symmetrical, in a high-volatility market, all options, whether calls or puts, higher or lower exercise prices, take on greater value. For the same reason, in a low-volatility market, all options take on reduced values.

Mean and Standard Deviation

If we assume a normal distribution of prices, we will need a method of describing the appropriate normal distribution to the theoretical pricing model. Fortunately, all normal distributions can be fully described with two numbers—the *mean* and the *standard deviation*. If we know that a distribution is normal, and we also know the mean and standard deviation, then we know all the characteristics of the distribution.

Graphically, we can interpret the mean as the location of the peak of the distribution and the standard deviation as a measure of how fast the distribution spreads out. Distributions that spread out very quickly, such as the one in Figure 6-4, have a high standard deviation. Distributions that spread out very slowly, such as the one in Figure 6-3, have a low standard deviation.

Numerically, the mean is simply the average outcome, a concept familiar to most traders. To calculate the mean, we add up all the results and divide by the total number of occurrences. Calculation of the standard deviation is not quite so simple and will be discussed later. What is important at this point is the interpretation of these numbers, in particular, what a mean and standard deviation suggest in terms of likely price movement.

Let's go back to Figure 6-2 and consider the troughs numbered 0 to 15 at the bottom. We suggested that these numbers might represent the number of heads resulting from 15 flips of a coin. Alternatively, they might represent the number of times a ball goes to the right at each nail as it drops down through the maze. The first trough is assigned 0 because any ball that ends there must go left at every nail. The last trough is assigned 15 because any ball that ends there must go right at every nail.

Suppose that we are told that the mean and standard deviation in Figure 6-2 are 7.50 and 3.00, respectively.[2] What does this tell us about the distribution? The mean tells us the average outcome. If we add up all the outcomes and divide by the number of occurrences, the result will be 7.50. In terms of the troughs, the average result will fall halfway between troughs 7 and 8. (Of course, this is not an actual possibility. However, we noted in Chapter 5 that the average outcome does not have to be an actual possibility for any one outcome.)

The standard deviation determines not only how fast the distribution spreads out, but it also tells us something about the likelihood of a ball ending up in a specific trough or group of troughs. In particular, the standard deviation tells us the probability of a ball ending up in a trough that is a specified distance from the mean. For example, we may want to know the likelihood of a ball falling down through the maze and ending up in a trough lower than 5 or higher than 10. The answer to this question depends on the number of standard deviations the ball must move away from the mean. If we know this, we can determine the probability associated with that number of standard deviations.

The exact probability associated with any specific number of standard deviations can be found in most texts on statistics or probability. Alternatively, such probabilities can be easily calculated in most commonly used computer spreadsheet programs. For option traders, the following approximations will be useful:

> ±1 standard deviation takes in approximately 68.3 percent (about 2/3) of all occurrences.

[2] The reader who is familiar with the mean and standard deviation and who would like to check the arithmetic will find that the actual mean and standard deviation are 7.49 and 3.02. For simplicity, we have rounded these to 7.50 and 3.00.

±2 standard deviations takes in approximately 95.4 percent (about 19/20) of all occurrences.

±3 standard deviations takes in approximately 99.7 percent (about 369/370) of all occurrences.

Note that each number of standard deviations is preceded by a plus or minus sign. Because normal distributions are symmetrical, the likelihood of up movement and down movement is identical. The probability associated with each number of standard deviations is usually given as a decimal value, but a fractional approximation is often useful to traders, and this appears in parentheses.

Now let's try to answer our question about the likelihood of getting a ball in a trough lower than 5 or higher than 10. We can designate the divider between troughs 7 and 8 as the mean, 7½. If the standard deviation is 3, what troughs are within one standard deviation of the mean? One standard deviation from the mean is 7½ ± 3, or 4½ to10½. Interpreting ½ as the divider between troughs, we can see that troughs 5 through 10 fall within 1 standard deviation of the mean. We know that one standard deviation takes in about two-thirds of all occurrences, so we can conclude that out of every three balls we drop into the maze, two should end up in troughs 5 through 10. Whatever is left over, one out of every three balls, will end up in one of the remaining troughs, 0 through 4 and 11 through 15. Hence, the answer to our original question about the likelihood of getting a ball in a trough lower than 5 or higher than 10 is about 1 chance in 3, or about 33 percent. (The exact answer is 100% – 68.3% = 31.7%.) This is shown in Figure 6-6.

Let's try another calculation, but this time we can think of the problem as a wager. Suppose that someone offers us 30 to 1 odds against dropping a ball into the maze and having it end up specifically in troughs 14 or 15. Is this bet worth making? One characteristic of standard deviations is that they are additive. In our example, if one standard deviation is 3, then two standard deviations are 6. Two standard deviations from the mean is therefore 7½ ± 6, or 1½ to 13½. We can see in Figure 6-6 that troughs 14 and 15 lie outside two standard deviations. Because the probability of getting a result within two standard deviations is approximately 19 out of 20, the probability of getting a result beyond two standard deviations is 1 chance in 20. Therefore, 30 to 1 odds may seem very favorable. Recall, however, that beyond two standard deviations also includes troughs 0 and 1. Because normal distributions are symmetrical, the chances of getting a ball specifically in troughs 14 or 15 must be half of 1 chance in 20, or about 1 chance in 40. At 30 to 1 odds, the bet must be a bad one because the odds do not sufficiently compensate us for the risk involved.

In Chapter 5, we suggested that a truly accurate theoretical pricing model would require us to assign probabilities to an infinite number of possible price outcomes for an underlying contract. Then, if we multiply each price outcome by its associated probability, we can use the results to calculate an option's theoretical value. The problem is that an infinite number of anything is not easy to work with. Fortunately, the characteristics of normal distributions are so well known that formulas have been developed that facilitate the computation of both the probabilities associated with every point along a normal distribution

Figure 6-6

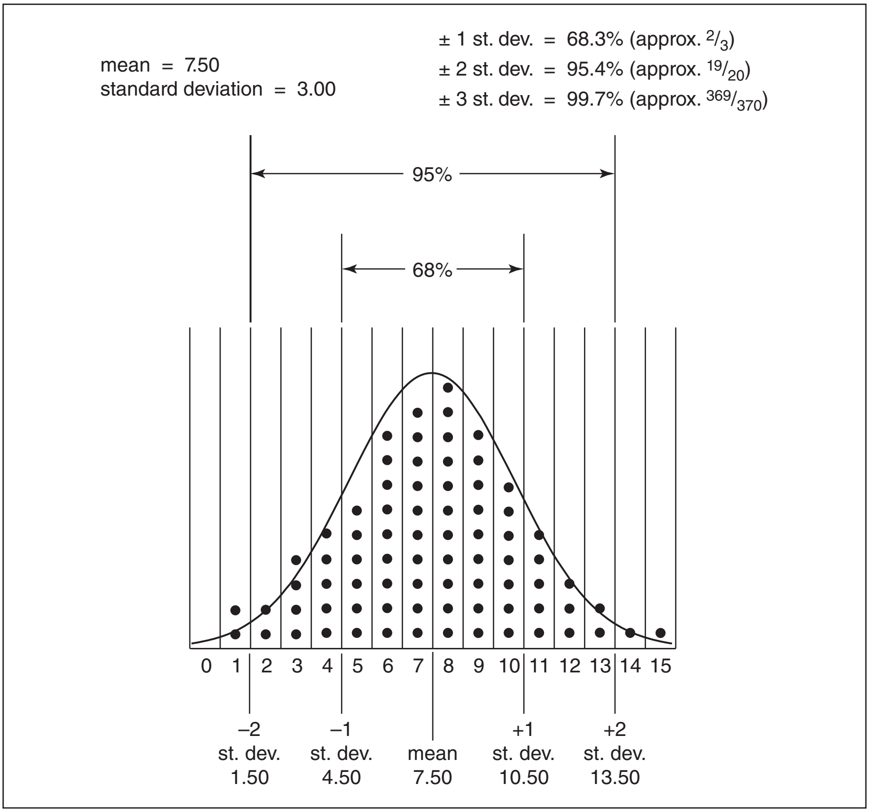

curve and the area under various portions of the curve. If we assume that prices of an underlying instrument are normally distributed, these formulas represent a unique set of tools to help us solve for an option's theoretical value.

Louis Bachelier was the first to make the assumption that the prices of an underlying contract are normally distributed. As we shall see, there are logical problems with this assumption. Consequently, over the years, the assumption has been modified to make it more consistent with real-world conditions. In its modified form, it is the basis for many theoretical pricing models, including the Black-Scholes model.

Forward Price as the Mean of a Distribution

If we decide to assign probabilities that are consistent with a normal distribution, how do we feed this distribution into a theoretical pricing model? Because all normal distributions can be described by a mean and a standard deviation, in some way we must feed these two numbers into our pricing model.

In Chapter 5, we suggested that any distribution ought to be centered around the most likely underlying price at expiration. Although we cannot know exactly what that price will be, if we assume that no arbitrage opportunity exists in the underlying contract, a logical guess is the forward price. If we make the assumption that the forward price represents the mean of a distribution, then in the long run, any trade made at the current underlying price will just break even. The various forms of the Black-Scholes model differ primarily in how they calculate the forward price. Depending on the type of underlying contract, whether a stock, a futures contract, or a foreign currency, the model takes the current underlying price, the time to expiration, interest rates, and, in the case of stocks, dividends to calculate the forward price. It then makes this the mean of the distribution.

Volatility as a Standard Deviation

In addition to the mean, to fully describe a normal distribution, we also need a standard deviation. When we input a volatility into a theoretical pricing model, we are actually feeding in a standard deviation. Volatility is just a trader's term for standard deviation. Because the Greek letter sigma (σ) is the traditional notation for standard deviation, in this text we will use the same notation for volatility.

At this point, it will help if we assign a working definition to volatility, although we will later modify this definition slightly. For the present, we will assume that the volatility we feed into a pricing model represents a one standard deviation price change, in percent, over a one-year period. For example, consider a contract with a one-year forward price of 100 and that we are told has a volatility of 20 percent. (We'll discuss later where this number might come from.) With a mean of 100 and a standard deviation of 20 percent, if we come back one year from now, there is a 68 percent probability that the contract will be trading between 80 and 120 (100 ± 20%), a 95 percent probability that the contract will be trading between 60 and 140 (100 ± 2 × 20%), and a 99.7 percent probability that the contract will be trading between 40 and 160 (100 ± 3 × 20%). These are the probabilities associated with one, two, and three standard deviations.

Instead of specifying the forward price, suppose that we are dealing with a stock that is currently trading at $100 and that has the same 20 percent volatility. In order to determine the one-year probabilities, we must first determine the one-year forward price because this represents the mean of the distribution. If interest rates are 8 percent and the stock pays no dividends, the one-year forward price will be $108. Now, a one standard deviation price change is 20% × $108 = $21.60. Thus, one year from now, we would expect the same stock to be trading between $86.40 and $129.60 ($108 ± $21.60) approximately 68 percent of the time, between $64.80 and $151.20 ($108 ± 2 × $21.60) approximately 95 percent of the time, and between $43.20 and $172.80 ($108 ± 3 × $21.60) approximately 99.7 percent of the time.

Returning to our contract with a forward price of 100, suppose that we come back at the end of one year and find that the contract, which we thought had a volatility of 20 percent, is trading at 35. Does this mean that the volatility of 20 percent was wrong? A price change of more than three standard deviations may be unlikely, but one should not confuse unlikely with impossible.

Flipping a perfectly balanced coin 15 times may result in 15 heads, even though the odds of this occurring are less than one chance in 32,000. If 20 percent is the right volatility, the probability that the price will fall from 100 to 35 in one year is less than one chance in 1,500. However, one chance in 1,500 is not impossible, and perhaps this was the one time in 1,500 when the price did indeed end up at 35. Of course, it is also possible that we had the wrong volatility. But we can't make that determination without looking at a large number of price changes for the contract so that we have a representative price distribution.

Scaling Volatility for Time

Like interest rates, volatility is always expressed as an annualized number. If someone says that interest rates are 6 percent, no one needs to ask whether that means 6 percent per day, 6 percent per week, or 6 percent per month. Everyone knows that it means 6 percent per year. The same is true of volatility.

We might logically ask what an annual volatility tells us about the likelihood of price changes over some shorter period of time. Although interest rates are proportional to time (we simply multiply the rate by the amount of time), volatility is proportional to the square root of time. To calculate a volatility, or standard deviation, over some period of time other than one year, we must multiply the annual volatility be the square root of time, where the time period t is expressed in years

$$\text{Volatility}_t = \text{volatility}_{\text{annual}} \times \sqrt{t}$$

Traders typically calculate volatility for an underlying contract by observing price changes at regular intervals. Let's begin by assuming that we plan to observe price changes at the end of every day. Because there are 365 days in a year, it might seem that prices can change 365 times per year. In this text, though, we are focusing primarily on exchange-traded contracts. Because most exchanges are closed on weekends and holidays, if we observe the price of an underlying contract at the end of every day, prices cannot really change 365 times per year. Depending on the exchange, there are probably somewhere between 250 and 260 trading days in a year.[3] Because we need the square root of the number of trading days, for convenience, many traders assume that there are 256 trading days in a year given that the square root of 256 is a whole number, 16. If we make this assumption, then

$$\text{Volatility}_{\text{daily}} = \text{volatility}_{\text{annual}} \times \sqrt{1/256} = \text{volatility}_{\text{annual}} \times 1/16 = \frac{\text{volatility}_{\text{annual}}}{16}$$

To approximate a daily standard deviation, we can divide the annual volatility by 16.

Returning to our contract trading at 100 with a volatility of 20 percent, what is a one standard deviation price change from one day to the next?

[3] As markets around the world become more integrated, and with the advent of electronic trading, it may become more difficult to determine exactly what fraction of a year one day represents. Depending on the contract and exchange, in some cases it may be sensible to look at prices every day, 365 days per year.

The answer is 20%/16 = 1¼%, so a one standard deviation daily price change is 1¼% × 100 = 1.25. We expect to see a price change of 1.25 or less approximately two trading days out of every three and a price change of 2.50 or less approximately 19 trading days out of every 20. Only one day in 20 would we expect to see a price change of more than 2.50.

We can do the same type of calculation for a weekly standard deviation. Now we must ask how many times per year prices can change if we look at prices once a week. There are no complete weeks when no trading takes place, so we must make our calculations using all 52 trading weeks in a year. Therefore,

$$\text{Volatility}_{\text{weekly}} = \text{volatility}_{\text{annual}} \times \sqrt{1/52} \approx \text{volatility}_{\text{annual}} \times 1/7.2 = \frac{\text{volatility}_{\text{annual}}}{7.2}$$

To approximate a weekly standard deviation, we can divide the annual volatility by 7.2. Dividing our annual volatility of 20 percent by the square root of 52, or approximately 7.2, we get 20%/7.2 » 2.78. For our contract trading at 100, we would expect to see a price change of 2.78 or less two weeks out of every three, a price change of 5.56 or less 19 weeks out of every 20, and only one week in 20 would we expect to see a price change of more than 5.56.

If we want to be as accurate as possible, when estimating a daily or weekly standard deviation, we ought to begin by calculating the one-day or one-week forward price. But for short periods of time, the forward price is so close to the current price that most traders assume for convenience that a one-day or one-week distribution is centered around the current price.

Suppose that a stock is trading at \$45 per share and has an annual volatility of 37 percent. What is an approximate one and two standard deviation price range from one day to the next and from one week to the next? For one day, we can divide the annual volatility by 16 (the square root of 256, the number of trading days in a year)

$$\$45 \times \frac{37\%}{16} \approx \$1.04$$

A one and two standard deviation daily price range is approximately

$$\$45 \pm \$1.04 \approx \$43.96 \text{ to } \$46.04 \quad \text{(one standard deviation)}$$
$$\$45 \pm (2 \times \$1.04) \approx \$42.92 \text{ to } \$47.08 \quad \text{(two standard deviations)}$$

For one week, we can divide the annual volatility by 7.2 (the square root of 52, the number of trading weeks in a year)

$$\$45 \times \frac{37\%}{7.2} = \$2.31$$

A one and two standard deviation weekly price range is approximately

$$\$45 \pm \$2.31 \approx \$42.69 \text{ to } \$47.31 \quad \text{(one standard deviation)}$$
$$\$45 \pm (2 \times \$2.31) \approx \$40.38 \text{ to } \$49.62 \quad \text{(two standard deviations)}$$

When we scale volatility for time, the same probabilities still apply. Approximately 68 percent of the occurrences will fall within one standard deviation. Approximately 95 percent of the occurrences will fall within two standard deviations.

Volatility and Observed Price Changes

Why might a trader want to estimate daily or weekly price changes from an annual volatility? Volatility is the one input into a theoretical pricing model that cannot be directly observed. Yet many option strategies, if they are to be successful, require a reasonable assessment of volatility. Therefore, an option trader needs some method of determining whether his expectations about volatility are being realized in the marketplace. Unlike directional strategies, whose success or failure can be immediately observed from current prices, there is no such thing as a current volatility. A trader must usually determine for himself whether he is using a reasonable volatility input into the theoretical pricing model.

Previously, we estimated that for a \$45 stock with an annual volatility of 37 percent, a one standard deviation price change is approximately \$1.04. Suppose that over five days we observe the following daily settlement price changes:

+\$0.98, –\$0.65, –\$0.70, +\$0.25, –\$0.85

Are these price changes consistent with a 37 percent volatility?

We expect to see a price change of more than \$1.04 (one standard deviation) about one day in three. Over five days, we would expect to see at least one day, and perhaps two days, with a price change greater than one standard deviation. Yet, during this five-day period, we did not see a price change greater than \$1.04 even once. What conclusions can be drawn from this? One thing seems clear: these five price changes do not appear to be consistent with a 37 percent volatility.

Before making any decisions, we ought to consider any unusual conditions that might be affecting the observed price changes. Perhaps this was a holiday week, and as such, it did not reflect normal market activity. If this is our conclusion, then 37 percent may still be a reasonable volatility estimate. On the other hand, if we can see no logical reason for the market being less volatile than predicted by a 37 percent volatility, then we may simply be using the wrong volatility. If we come to this conclusion, perhaps we ought to consider using a lower volatility that is more consistent with the observed price changes. If we continue to use a volatility that is not consistent with the actual price changes, then we have the wrong volatility. If we have the wrong volatility, we have the wrong probabilities. And if we have the wrong probabilities, we are generating incorrect theoretical values, thereby defeating the purpose of using a theoretical pricing model in the first place.

Admittedly, five days is a very small number of price changes, and it is unlikely that a trader will rely heavily on such a small sample. If we flip a coin five times and it comes up heads each time, we may not be able to draw any definitive conclusions. But if we flip the coin 50 times and it comes up heads every time, now we might conclude that there is something wrong with the coin. In the same way, most traders prefer to see larger price samplings, perhaps 20 days, or 50 days, or 100 days, before drawing any dramatic conclusions about volatility.

Exactly what volatility is associated with the five price changes in the foregoing example? Without doing some rather involved arithmetic, it is difficult to say.

(The answer is actually 27.8 percent.) However, if a trader has some idea of the price changes he expects, he can easily see that the changes over the five-day period are not consistent with a 37 percent volatility.[4]

We have used the phrase *price change* in conjunction with our volatility estimates. Exactly what do we mean by this? Do we mean the high/low during some period? Do we mean open-to-close price changes? Or is there another interpretation? Although various methods have been suggested to estimate volatility, the most common method for exchange-traded contracts has been to calculate volatility based on settlement-to-settlement price changes. Using this approach, when we say that a one standard deviation daily price change is $1.04, we mean $1.04 from one day's settlement price to the next day's settlement price. The high/low or open/close price range may have been either more or less than this amount, but it is the settlement-to-settlement price change on which we focus.[5]

A Note on Interest-Rate Products

For some interest-rate products, primarily Eurocurrency interest-rate futures, the listed contract price represents the interest rate associated with that contract, expressed as a whole number, subtracted from 100.[6] If the London Interbank Offered Rate (LIBOR), the interest paid on dollar deposits outside the United States, is 7.00 percent, the associated Eurodollar futures contract traded at the Chicago Mercantile Exchange will be trading at 100 – 7.00 = 93.00. If Euro Interbank Offered Rate (Euribor), the interest paid on euro deposits outside the European Economic Union, is 4.50 percent, the associated Euribor futures contract traded at the London International Financial Futures Exchange will be trading at 100 – 4.50 = 95.50. Volatility calculations for these contracts are done using the rate associated with the contract (the *rate volatility*) rather than the price of the contract (the *price volatility*).

If a Eurodollar futures contract is trading at 93.00 with a volatility of 26 percent, an approximate daily and weekly one standard deviation price change is

$$(100 - 93) \times \frac{26\%}{16} \approx 0.11 \quad \text{(daily standard deviation)}$$

$$(100 - 93) \times \frac{26\%}{7.2} \approx 0.25 \quad \text{(weekly standard deviation)}$$

---

[4] A price change greater than two standard deviations will occur about 1 time in 20. Because there are approximately 20 trading days in a month, as an additional benchmark, most traders expect to see a daily two standard deviation occurrence about once a month.

[5] Alternative methods of estimating volatility have also been proposed when trading is continuous or when there is no well-defined daily settlement price. See, for example, Michael Parkinson, "The Extreme Value Method of Estimating the Variance of the Rate of Return," *Journal of Business* 53(1):61–64, 1980; Mark B. Garman and Michael J. Klass, "On the Estimation of Security Price Volatilities from Historical Data," *Journal of Business* 53(1):67–78, 1980; and Stan Beckers, "Variance of Security Price Returns Based on High, Low, and Closing Prices," *Journal of Business* 56(1):97–112, 1983.

[6] This method of quoting Eurocurrency contracts is used so that moves in Eurocurrency contracts will tend to mimic moves in bond prices. If interest rates rise, both bond prices and Eurocurrency futures will fall; if interest rates fall, both bond prices and Eurocurrency futures will rise.

To be consistent, if we index Eurodollar futures prices from 100, we must also index exercise prices from 100. Therefore, a 93.00 exercise price in our pricing model is really a 7.00 percent (100 – 93.00) exercise price. This transformation also requires us to reverse the type of option, changing calls to puts and puts to calls. To see why, consider a 93.00 call. For this call to go into the money, the underlying contract must rise above 93.00. But this requires that interest rates fall below 7.00 percent. Therefore, a 93.00 call in listed terms is the same as a 7.00 percent put in interest-rate terms. A model that is correctly set up to evaluate options on Eurodollar or other types of indexed interest-rate contracts will make this transformation automatically. The price of the underlying contract and the exercise price are subtracted from 100, with listed calls treated as puts and listed puts treated as calls.

This type of transformation is not required for most bonds and notes. Depending on the coupon rate, the prices of these products may range freely without upper limit, often exceeding 100. Exchange-traded options on bond and note futures are therefore most often evaluated using a traditional pricing model. However, interest-rate products present other problems that may require specialized pricing models.

It is possible to take an instrument such as a bond and calculate the current yield based on its price in the marketplace. If we were to take a series of bond prices and from these calculate a series of yields, we could calculate the *yield volatility*, that is, the volatility based on the change in yield. We might then use this number to evaluate the theoretical value of an option on the bond, although to be consistent we would also have to specify the exercise price in terms of yield. Because it is possible to calculate the volatility of an interest-rate product using these two different methods, interest-rate traders usually make a distinction between yield volatility (the volatility calculated from the current yield on the instrument) and price volatility (the volatility calculated from the price of the instrument in the marketplace).

Lognormal Distributions

Thus far we have assumed that the prices of an underlying instrument are normally distributed. Is this a reasonable assumption? Beyond the question of the exact distribution of prices in the real world, the normal distribution assumption has one serious flaw. A normal distribution is symmetrical. For every possible upward move in the price of an underlying instrument, there must be the possibility of a downward move of equal magnitude. If we allow for the possibility of a $50 contract rising $75 to $125, we also must allow for the possibility of the contract dropping $75 to a price of –$25. But negative prices are clearly not possible for traditional stocks or commodities.

We have defined volatility in terms of the percent changes in the price of an underlying instrument. In this sense, an interest rate and volatility are similar in that they both represent a *rate of return*. The primary difference between interest and volatility is that interest accrues only at a positive rate, whereas volatility is a combination of positive and negative rates of return. If we invest money at a fixed interest rate, the value of the principal will always grow.

However, if we invest in an underlying instrument with a volatility other than zero, the instrument may go up or down in price, resulting in either a profit (a positive rate of return) or a loss (a negative rate of return).

A rate-of-return calculation must specify not only the rate that is being used but also the time intervals over which the returns are calculated. Suppose that we invest $1,000 for one year at an annual interest rate of 12 percent. How much will we have at the end of one year? The answer depends on how the 12 percent interest on our investment is paid out.

Rate of Payment	Value after One Year	Total Yield
12% once a year	$1,120.00	12.00%
6% twice a year	$1,123.60	12.36%
3% every three months	$1,125.51	12.55%
1% every month	$1,126.83	12.68%
12%/52 every week	$1,127.34	12.73%
12%/365 every day	$1,127.47	12.75%
12% compounded continuously	$1,127.50	12.75%

As interest is paid more frequently, even though it is paid at the same rate of 12 percent per year, the total yield on the investment increases. The yield is greatest when interest is paid continuously. In this case, it is just as if interest is paid at every possible moment in time.

Although less common, we can do the same type of calculation using a negative interest rate. For example, suppose that we make a bad investment of $1,000 and lose money at a rate of 12 percent annually (interest rate = –12%). How much will we have at the end of a year? The answer, again, depends on the frequency at which our losses accrue.

Rate of Payment	Value after One Year	Total Yield
–12% once a year	$880.00	–12.00%
–6% twice a year	$883.60	–11.64%
–3% every three months	$885.29	–11.47%
–1% every month	$886.38	–11.36%
–12%/52 every week	$886.80	–11.32%
–12%/365 every day	$886.90	–11.31%
–12% compounded continuously	$886.92	–11.31%

In the case of a negative interest rate, as losses are compounded more frequently, even though at the same rate of –12 percent per year, the smaller the total loss, and consequently, the smaller the negative yield.

In the same way that interest can be compounded at different intervals, volatility can also be compounded at different intervals. The Black-Scholes model is a *continuous-time* model. The model assumes that volatility is compounded continuously, just as if the price changes in the underlying contract, either up or down, are taking place continuously but at an annual rate corresponding to the volatility associated with the contract. When the percent price changes are normally distributed, the continuous compounding of these price changes will result in a *lognormal distribution* of prices at expiration. Such a distribution is shown in Figure 6-7. The entire distribution is skewed toward the upside because upside price changes (a positive rate of return) will be greater, in absolute terms, than downside price changes (a negative rate of return). In our interest-rate example, a continuously compounded rate of return of +12 percent yields a profit of $127.50 after one year, whereas a continuously compounded rate of return of –12 percent yields a loss of only $113.08. If the 12 percent is a volatility, then a one standard deviation upward price change at the end of one year is +$127.50, whereas a one standard deviation downward price change is –$113.08. Even though the rate of return is a constant 12 percent, the continuous compounding of 12 percent yields different upward and downward moves.

Note also the location of the mean of the distributions in Figure 6-7. The mean can be thought of as the “balance point” of the distribution. For a normal distribution, the peak of the distribution, or *mode*, and the mean have the same location, exactly in the middle of the distribution. But in a lognormal distribution the right tail, which is open-ended, is longer than the left tail, which is bounded by zero. Because there is more “weight” to the right of the peak, the mean of the lognormal distribution must be located to the right of the peak.

Figure 6-7

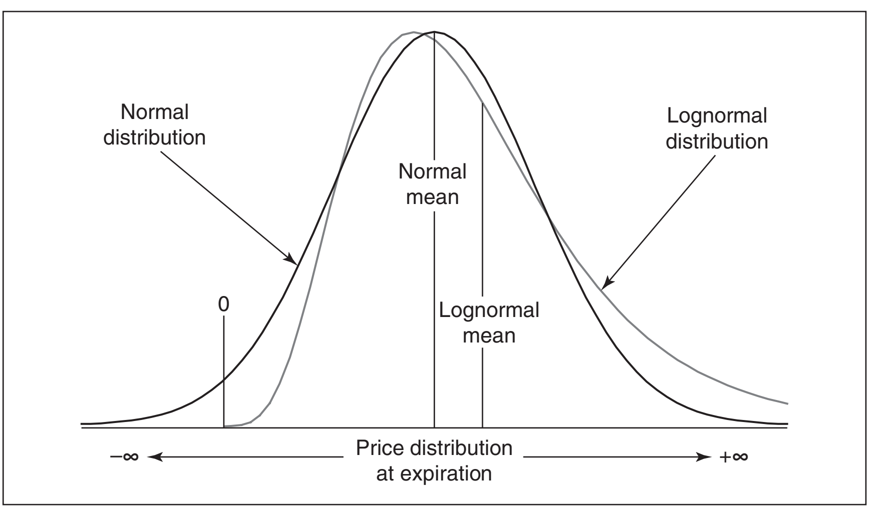

Continuous rates of return can be calculated using the exponential function,[7] denoted by either exp(x) or e^x. In the preceding examples,

$$\$1{,}000 \times e^{0.12} = \$1{,}127.50 \quad \text{and} \quad \$1{,}000 \times e^{-0.12} = \$886.92$$

No matter how large the negative interest rate, continuous compounding precludes the possibility of an investment falling below zero because it is impossible to lose more than 100 percent of an investment. Consequently, in a lognormal distribution, the value of the underlying instrument is bounded by zero on the downside. Clearly, this is a more realistic representation of the real world than a normal distribution.

We can see the effect of using a lognormal distribution rather than a normal distribution by considering the value of a 90 put and 110 call with a forward price of 100 for the underlying contract with six months to expiration and a volatility of 30 percent

Contract	Value if Price Distribution Is Normal	Value if Price Distribution Is Lognormal
90 put	4.37	4.00
110 call	4.37	4.74

Under a normal distribution assumption, both the call and put have exactly the same value because they are both 10 percent out of the money. But under the lognormal distribution assumption in the Black-Scholes model, the 110 call will always have a greater value than the 90 put. The value of the 110 call can potentially appreciate without limit because the price of the underlying contract has no limit on the upside. The 90 put, however, can only rise to a maximum value of 90 because the price of the underlying contract can never fall below zero.

Of course, the values in the preceding example are true only in theory. There is no law that prevents the 90 put from trading at a price greater than the 110 call. Indeed, such price relationships occur in many markets for a variety of reasons that we will discuss later. However, one possible explanation is that the marketplace disagrees with the assumptions on which the model is based. Perhaps the marketplace believes that a lognormal distribution is not an accurate representation of possible prices. And perhaps the marketplace is right!

Interpreting Volatility Data

When traders discuss volatility, even experienced traders may find that they are not always talking about the same thing. When a trader says that the volatility is 25 percent, this statement may take on a variety of meanings. We can avoid confusion in subsequent discussions if we define some of the different ways in

[7] It will be useful for an option trader to become familiar with the characteristics of the exponential function [e^x or exp(x)] and its inverse, the logarithmic function [ln(x)]. These can be found in any algebra or finance text.

which traders refer to volatility. We can begin by dividing volatility into two categories—*realized volatility*, which we associate with an underlying contract, and *implied volatility*, which we associate with options.

Realized Volatility

The realized volatility is the annualized standard deviation of percent price changes of an underlying contract over some period of time.[8] When we calculate realized volatility, we must specify both the interval at which we are measuring the price changes and the number of intervals to be used in the calculations. For example, we might talk about the 50-day volatility of an underlying contract. Or we might talk about the 52-week volatility of a contract. In the former case, we are calculating the volatility from the daily price changes over a 50-day period.[9] In the latter case, we are calculating the volatility from the weekly price changes over a 52-week period.

On a graph of realized volatility, each point represents the volatility over a specified period using price changes over a specified interval. If we graph the 50-day volatility of a contract, each point on the graph represents the annualized standard deviation of the daily price changes over the previous 50 days. If we graph the 52-week volatility, each point on the graph represents the annualized standard deviation of the weekly price changes over the previous 52 weeks.

Traders may also refer to realized volatility in the future (*future realized volatility*) and realized volatility in the past (*historical realized volatility*). The future realized volatility is what every trader would like to know—the volatility that best describes the future distribution of price changes for an underlying contract. In theory, it is the future realized volatility over the life of the option that we need to input into a theoretical pricing model. If a trader knows the future realized volatility, he knows the right "odds." When he feeds this number into a theoretical pricing model, he can generate accurate theoretical values because he has the right probabilities. Like the casino, he may lose in the short run because of bad luck, but in the long run, with the probabilities in his favor, the trader can be reasonably certain of making a profit.

Clearly, no one knows what the future holds. However, if a trader intends to use a theoretical pricing model, he must try to make an estimate of future realized volatility. In option evaluation, as in other disciplines, a good starting point is historical data. What typically has been the historical realized volatility of a contract? If, over the past 10 years, the volatility of a contract has never been less than 10 percent nor more than 30 percent, a guess for the future volatility of either 5 or 40 percent hardly makes sense. This does not mean that either of these extremes is impossible. But based on past performance, and in the absence of any extraordinary circumstances, a guess within the historical limits of 10 and 30 percent is probably more realistic than a guess outside these limits. Of course,

[8] In order to turn price changes into continuously compounded returns, volatility is most often calculated using logarithmic price changes—the natural logarithm of the current price divided by the previous price. In most cases, there is little practical difference between the percent price changes and logarithmic price changes.

[9] For exchange-traded contracts, volatility calculations using daily intervals typically include only business days because these are the only days on which prices can actually change. If there are five trading days per week, a 50-day volatility covers a period of approximately 10 weeks.

10 to 30 percent is still a very wide range. But at least the historical data offers a starting point. Additional information may help to further narrow the estimate.

As option traders have come to appreciate the importance of volatility as an input into a pricing model, volatility forecasting models have been developed in an attempt to more accurately predict future realized volatility. If a trader has access to a volatility *forecast* that he believes is reliable, he will want to use this forecast to make a better decision as to the future realized volatility. We will put off a discussion of possible forecasting methods until later chapters.

When we calculate volatility over a given period of time, we still have a choice of the time intervals over which to measure the price changes in the underlying contract. A trader might consider whether the choice of intervals, even if the intervals cover the same time period, might affect the results. For example, we might look at the 250-day volatility, the 52-week volatility, and the 12-month volatility of a contract. All volatilities cover approximately one year, but one is calculated from daily price changes, one from weekly price changes, and one from monthly price changes.

For most underlying contracts, the interval that is chosen does not seem to greatly affect the result. It is possible that a contract will make large daily moves yet finish the week unchanged. However, this is by far the exception. A contract that is volatile from day to day is likely to be equally volatile from week to week or month to month. Figure 6-8 shows the 250-day realized volatility of the S&P 500 Index from 2003 through 2012, with the volatility calculated from price changes at three different intervals: daily, weekly, and every four weeks. The graphs are not identical, but they do seem to have similar characteristics. There is no clear evidence that using one interval rather than another results in consistently higher or lower volatility.

Figure 6-8 S&P 500 Index 250-day historical volatility.

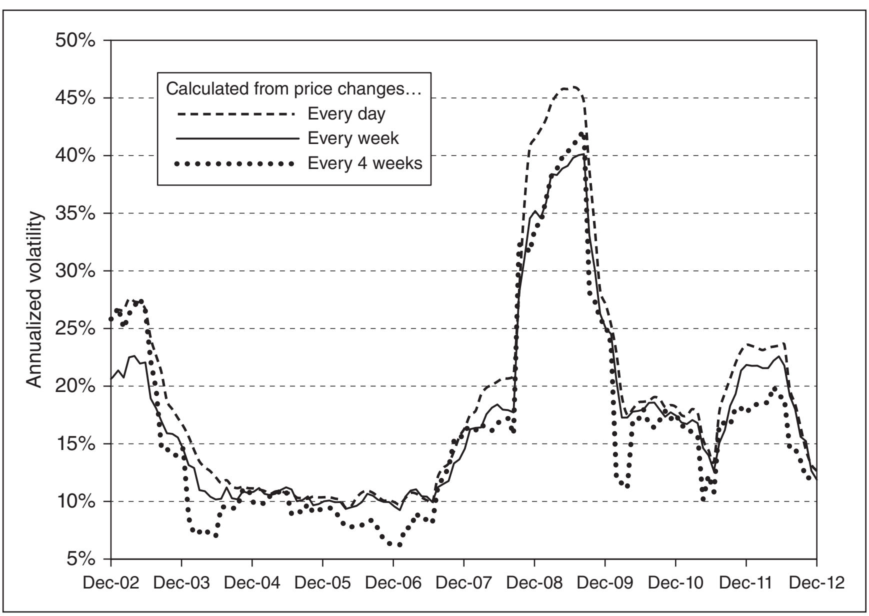

Implied Volatility

Unlike realized volatility, which is calculated from price changes in the underlying contract, implied volatility is derived from the price of an option in the marketplace. In a sense, the implied volatility represents the marketplace's consensus of what the future realized volatility of the underlying contract will be over the life of the option.

Consider a three-month 105 call on a stock that pays no dividend. If we are interested in purchasing this call, we might use a pricing model to determine the option's theoretical value. For simplicity, let's assume that the option is European (no early exercise) and that we will use the Black-Scholes model. In addition to the exercise price, time to expiration, and type, we also need the price of the stock, an interest rate, and a volatility. Suppose that the current stock price is 98.50, the three-month interest rate is 6.00 percent, and our best estimate of volatility over the next three months is 25 percent. When we feed this data into our model, we find that the option has a theoretical value of 2.94. However, when we check the price of the option in the marketplace, we find that the 105 call is trading very actively at a price of 3.60. How can we account for the fact that we think the option is worth 2.94, but the rest of the world seems to think that it's worth 3.60?

This is not an easy question to answer because there are many forces at work in the marketplace that cannot be easily identified or quantified. But one way we might try to answer the question is by making the assumption that everyone trading the option is using the same theoretical pricing model. If we make this assumption, the cause of the discrepancy must be a difference of opinion about one or more of the inputs into the model. Which inputs are the most likely cause?

It's unlikely to be either the time to expiration or the exercise price because these inputs are fixed in the option contract. What about the underlying price of 98.50? Perhaps we incorrectly estimated the stock price due to the width of the bid-ask spread. However, for most actively traded underlying contracts, it is unlikely that the spread will be wide enough to cause a discrepancy of 0.66 in the value of the option. In order to yield a value of 3.60 for the 105 call, we would actually have to raise the stock price to 100.16, and this is almost certainly well outside the bid-ask spread for the stock.

Perhaps our problem is the interest rate of 6.00 percent. But interest rates are usually the least important of the inputs into a theoretical pricing model. In fact, we would have to make a huge change in the interest-rate input, from 6.00 to 13.30 percent, to yield a theoretical value of 3.60.

This leaves us with one likely cause for the discrepancy—the volatility. In a sense, the marketplace seems be using a volatility that is different from 25 percent. To determine what volatility the marketplace is using, we can ask the following question: if we hold all other inputs constant (i.e., time to expiration, exercise price, underlying price, and interest rates), what volatility must we feed into our model to yield a theoretical value equal to the price of the option in the marketplace? In our example, we want to know what volatility will yield a value of 3.60 for the 105 call. Clearly, the volatility has to be higher than 25 percent, so we might begin to raise the volatility input into our model.

Figure 6-9

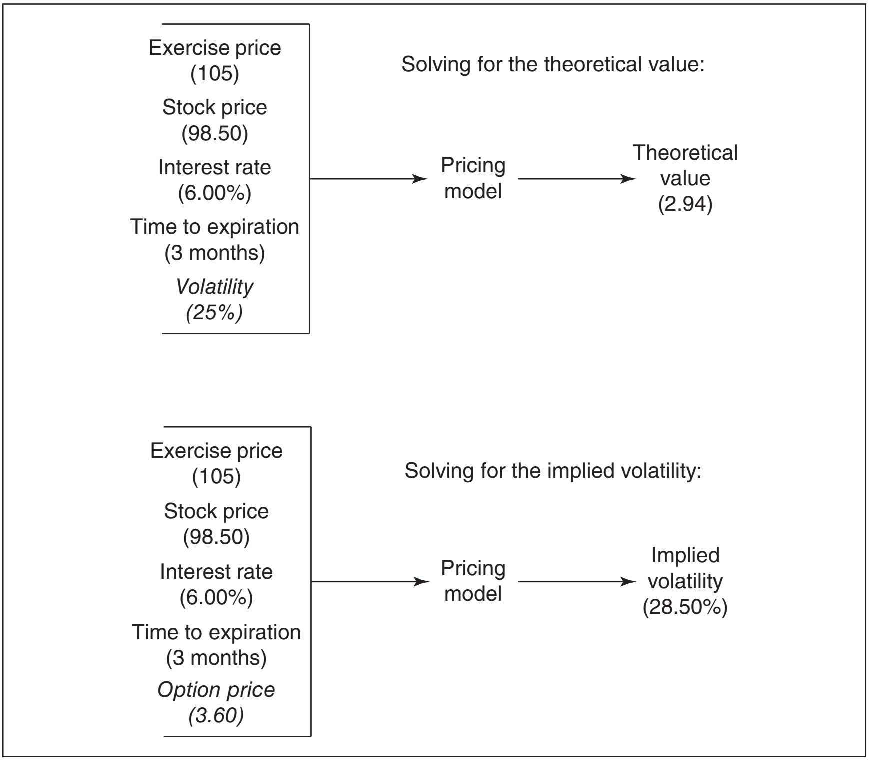

If we do, we find that at a volatility of 28.50 percent, the 105 call has a theoretical value of 3.60. The implied volatility of the 105 call—the volatility being implied to the underlying contract through the pricing of the option in the marketplace—is 28.50 percent.

When we solve for the implied volatility of an option, we are assuming that the theoretical value (the option's price), as well as all other inputs except volatility, are known. In effect, we are running the theoretical pricing model backwards to solve for the unknown volatility. In practice, this is easier said than done because most theoretical pricing models do not work in reverse. However, there are a number of relatively simple algorithms that can quickly solve for the implied volatility when all other inputs are known.

Implied volatility depends not only on the inputs into the theoretical pricing model but also on the theoretical pricing model being used. For some options, different models can yield significantly different implied volatilities. Problems can also arise when the inputs are not contemporaneous. If an option has not traded for some time and market conditions have changed, using an outdated option price will result in a misleading or inaccurate implied volatility. Suppose in our example that the price of 3.60 for the 105 call reflected the last trade, but that trade took place two hours ago when the underlying stock price was actually 99.25. At a stock price of 99.25, the implied volatility of the option, at a price of 3.60, is actually 26.95 percent. This underscores the importance of accurate and contemporaneous inputs when calculating implied volatilities.

Brokerage firms and data vendors who provide option analysis for their clients will typically include implied volatility data. The data may incorporate implied volatilities for every option on an underlying contract, or the data may be in the form of one implied volatility that is representative of options on a particular underlying market. In the latter case, the single implied volatility is usually the result of weighting the individual implied volatilities by some criteria, such as volume of options traded or open interest, or, as is most common, by assigning the greatest weight to the at-the-money options.

Implied volatility in the marketplace is constantly changing because option prices, as well as other market conditions, are constantly changing. It is as if the marketplace were continuously polling all the participants to come up with a consensus volatility for the underlying contract for each expiration. This is not a poll in the true sense because the traders do not confer with each other and then vote on the correct volatility. However, as bids and offers are made, the price at which an option is trading will represent an equilibrium between supply and demand. This equilibrium can be expressed as an implied volatility.

While the term *premium* really refers to an option's price, because the implied volatility is derived from an option's price, traders sometimes use premium and implied volatility interchangeably. If the current implied volatility is high by historical standards or high relative to the recent historical volatility of the underlying contract, a trader might say that premium levels are high; if implied volatility is unusually low, he might say that premium levels are low.

New option traders are taught, quite sensibly, to sell overpriced options and buy underpriced options. By selling options at prices higher than theoretical value or buying options at prices lower than theoretical value, a trader creates a positive theoretical edge. But how should a trader determine the degree to which an option is overpriced or underpriced? This sounds like an easy question to answer. Isn't the amount of the mispricing equal to the difference between the option's price and its value? The question arises because there is more than one way to measure this difference. Returning to our example of the 105 call, we might say that with a theoretical value of 2.94 and a price of 3.60, the 105 call is 0.66 overpriced. But in volatility terms the option is 3.50 volatility points overpriced because its theoretical value is based on a volatility of 25 percent (our volatility estimate), while its price is based on a volatility of 28.50 percent (the implied volatility). Given the unusual characteristics of options, it is often more useful for a trader to consider an option's price in terms of implied volatility rather than total points.

Implied volatility is often used by traders to compare the relative pricing of options. In our example, the 105 call is trading at 3.60 with an implied volatility of 28.50 percent. Suppose that a 100 call under the same conditions is trading at 5.40. In total points, the 100 call is clearly more expensive than the 105 call (5.40 versus 3.60). But if, at a price of 5.40, the 100 call has an implied volatility of 27.51 percent, most traders will conclude that in theoretical terms the 100 call is almost a full percentage point less expensive (27.51 percent versus 28.50 percent) than the 105 call. Traders, in fact, talk about implied volatility as if it were the price of an option. A trader who buys the 100 call at a price of 5.40 might say that she bought the call at 27.51 percent. A trader who sells the 105 call at a price of 3.60 might say that he sold the call at 28.50 percent. Of course, options are really bought and sold in the appropriate currency. But from an option trader's point of view the implied volatility is often a more useful expression of an option's price than its actual price in currency units.

Even if the implied volatility of the100 call is 27.51 percent and the implied volatility of the 105 call is 28.50 percent, this does not necessarily mean that a trader ought to buy the 100 call and sell the 105 call. A trader also will need to consider what will happen if his estimate of volatility turns out to be incorrect. If the future realized volatility over the life of the options turns out to be 25 percent, both the 100 call and the 105 call are overpriced, and the sale of either option should, in theory, result in a profit. But what will happen if the trader's volatility estimate is wrong, and the future realized volatility turns out to be 32 percent? Now the sale of either option will result in a loss. The consequences of being wrong about volatility are an important consideration, and this is something we will look at more closely in subsequent chapters. However, in the absence of other considerations, the lower implied volatility of the 100 call suggests that it is likely to be the better value.

Although option traders may at times refer to any of the various interpretations of volatility, two of these stand out in importance—the future realized volatility and the implied volatility. The future realized volatility of an underlying contract determines the *value* of options on that contract. The implied volatility is a reflection of an option's *price*. These two numbers, value and price, are what all traders, not just option traders, are concerned with. If a contract has a high value and a low price, a trader will want to be a buyer. If a contract has a low value and a high price, a trader will want to be a seller. For an option trader, this usually means comparing the expected future realized volatility with the implied volatility. If implied volatility is low with respect to the expected future volatility, a trader will prefer to buy options; if implied volatility is high, a trader will prefer to sell options. Of course, future volatility is an unknown, so a trader will look at historical and, if available, forecast volatility to help in making an intelligent guess about the future. In the final analysis, though, it is the future realized volatility that determines an option's value.

A commonly used analogy to help new traders better understand the role of volatility is to think of volatility as being similar to the weather. Suppose that a trader living in Chicago gets up on a July morning and must decide what clothes to wear that day. Will he consider putting on a heavy winter coat? This is probably not a logical choice because he knows that *historically* it is not sufficiently cold in Chicago in July to warrant wearing a winter coat. Next, he might turn on the radio or television to listen to the weather *forecast*. The forecaster is predicting clear skies with very warm temperatures close to 90°F (32°C). Based on this information, the trader has decided that he will wear a short-sleeve shirt and does not need a sweater or jacket. And he certainly won't need an umbrella. However, just to be sure, he decides to look out the window to see what the people passing in the street are wearing. To his surprise, everyone is wearing a coat and carrying an umbrella. Through their choice of clothing, the people outside are *implying* different weather than the forecast. Given the conflicting information, what clothes should the trader wear? He must make some decision, but whom should he believe, the weather forecaster or the people in the street? There can be no certain answer because the trader will not know the *future* weather until the end of the day. Much will depend on the trader's knowledge of local conditions. Perhaps the trader lives in an area far removed from where the weather forecaster is located. Then he must give added weight to local conditions.

The decision on what clothes to wear, like every trading decision, depends on a great many factors. Not only must the decision be made on the basis of the best available information, but the decision must also be made with consideration for the possibility of error. What are the benefits of being right? What are the consequences of being wrong? If a trader fails to take an umbrella and it rains, this may be of little consequence if the bus picks him up right outside his residence and drops him off right outside his place of work. On the other hand, if he must walk several blocks in the rain, he might become sick and have to miss several days of work. The choices are never easy, and one can only hope to make the decision that will turn out best in the long run.

Changing our assumptions about volatility can often have a dramatic effect on the value of an option. Figure 6-10 shows the prices, theoretical values, and implied volatilities for several gold options on July 31, 2012. Figure 6-11 focuses specifically on how these values change as we increase volatility from 14 to 18 percent. Looking for the moment at call values, although all the options increase in value, the 1600 call, the at-the-money option, increases the most, rising from 41.65 to 51.60, a total of 9.95. At the same time, the 1800 call shows the greatest increase in percent terms. Its value more than triples from 0.78 to 3.05, a total increase of 291 percent. These are important principles to which we will return later but that are worth stating now:

1. In total points, a change in volatility will have a greater effect on an at-the-money option than on an equivalent in-the-money or out-of-the-money option.
2. In percent terms, a change in volatility will have a greater effect on an out-of-the-money option than on an equivalent in-the-money or at-the-money option.

These same principles apply to puts as well as calls. The 1600 put increases the most in total points, rising from 29.26 to 39.21, a total of 9.95. The 1400 put increases the most in percent terms, from 0.13 to 0.89, or 585 percent.

No matter how one measures change, in-the-money options tend to be the least sensitive to changes in volatility. As an option moves deeply into the money, it becomes more sensitive to changes in the underlying price and less sensitive to changes in volatility. Because it is often volatility characteristics that investors and traders are looking for when they go into an options market, it should not come as a surprise that most of the trading volume in option markets is concentrated in at-the-money and out-of-the-money options, the options that are most sensitive to changes in volatility.

Figure 6-10 Gold eight-week (40 trading days) historical volatility.

July 31, 2012 October gold futures = 1612.4 Time to October expiration = 8 weeks (56 days) Interest rate = 0.50%*					
			Theoretical Value If Volatility Is ...		
Exercise Price	Settlement Price	Implied Volatility	Volatility = 14%	Volatility = 18%	Volatility = 22%
October calls					
1400	215.2	22.36%	212.37	213.13	214.98
1500	122.5	19.01%	116.05	121.01	127.34
1600	50.8	17.68%	41.65	51.6	61.57
1700	16.1	18.42%	8.08	15.28	23.49
1800	5.3	20.46%	0.78	3.05	7
October puts					
1400	2.9	22.26%	0.13	0.89	2.74
1500	10.2	19.02%	3.74	8.7	15.02
1600	38.4	17.68%	29.26	39.21	49.18
1700	103.7	18.45%	95.61	102.82	111.03
1800	192.8	20.50%	188.23	190.51	194.46

*The prices in Figures 6-10 and 6-12 occurred during a period of unusually low interest rates.

In Figures 6-12 and 6-13, we can see that the same principles apply to longer-term options. The at-the-money options (the December 1600 call and put) change most in total points, whereas the out-of-the-money options (the December 1800 call and 1400 put) change most in percent terms. As we would expect, the December option values are greater than the October option values with the same exercise price. But look at the magnitude of the changes as we change volatility. For the same exercise price, in total points, the December (long-term) options always change more than the October (short-term) options. This leads to a third principle of option evaluation:

3. A change in volatility will have a greater effect on a long-term option than an equivalent short-term option.

The reader may have noticed several interesting points in the foregoing figures. First, although implied volatilities may vary across exercise prices, calls and puts with the same exercise price and that expire at the same time have very similar implied volatilities. Second, when we change volatility, calls and

Figure 6-11

July 31, 2012
October gold futures = 1612.4
Time to October expiration = 8 weeks (56 days)
Interest rate = 0.50%

Exercise Price	Volatility = 14%	Volatility = 18%	Net Change in Value	Percent Change in Value
October calls				
1400	212.37	213.13	0.76	<1%
1500	116.05	121.01	4.96	4.00%
1600	41.65	51.6	9.95	24.00%
1700	8.08	15.28	7.2	89.00%
1800	0.78	3.05	2.27	291.00%
October puts				
1400	0.13	0.89	0.76	585.00%
1500	3.74	8.7	4.96	133.00%
1600	29.26	39.21	9.95	34.00%
1700	95.61	102.82	7.21	8.00%
1800	188.23	190.51	2.28	1.00%

puts with the same exercise price and time to expiration change by approximately the same amount. These characteristics are the result of an important relationship[10] between calls and puts at the same exercise price, a relationship that we will examine in more detail in Chapter 15.

Finally, we might ask how much the volatility of gold can change over an eight-week period? Is a 4 percentage point change a real possibility? In fact, from Figure 6-14, the eight-week historical volatility for the 3½ years leading up to July 2012, we can see that such changes are not at all uncommon.

Given its importance, it is not surprising that serious option traders spend a considerable amount of time thinking about volatility. From the historical, forecast, and implied volatility, a trader must try to make an intelligent decision about future volatility. From this, he will try to choose option strategies that will be profitable when he is right but that will not result in a serious loss when he is wrong. Because of the difficulty in predicting volatility, a trader must always look for strategies that will leave the greatest margin for error. No trader will survive very long pursuing strategies based on a future volatility estimate of 20 percent if such a strategy results in a significant

[10] Some readers may already be familiar with this relationship—*put-call parity*.

Figure 6-12

July 31, 2012
December gold futures = 1614.6
Time to December expiration = 17 weeks (119 days)
Interest rate = 0.50%

			Theoretical Value If Volatility Is ...		
Exercise Price	Settlement Price	Implied Volatility	Volatility = 14%	Volatility = 18%	Volatility = 22%
December calls					
1400	226.3	22.00%	216.03	220.06	226.3
1500	142.7	20.17%	126.41	136.59	148.05
1600	78.8	19.51%	58.78	73.31	87.86
1700	40	19.93%	20.71	33.52	47.1
1800	20.4	21.07%	5.44	13.03	22.84
December puts					
1400	11.9	21.92%	1.78	5.81	12.04
1500	28.2	20.14%	12	22.18	33.64
1600	64.2	19.50%	44.2	58.74	73.25
1700	125.3	19.94%	105.98	118.79	132.36
1800	205.5	21.07%	190.54	198.13	207.94

loss when volatility actually turns out to be 18 or 22 percent. Given the shifts that occur in volatility, a 2 percentage point margin for error may be no margin for error at all.

We have not yet concluded our discussion of volatility. But before continuing, it will be useful to look at option characteristics, trading strategies, and risk considerations. We will then be in a better position to examine volatility in greater detail.

Figure 6-13

July 31, 2012
December gold futures = 1614.6
Time to December expiration = 17 weeks (119 days)
Interest rate = 0.50%

Exercise Price	Volatility = 14%	Volatility = 18%	Net Change in Value	Percent Change in Value
December calls				
1400	216.03	220.06	4.03	2%
1500	126.41	136.59	10.18	8%
1600	58.78	73.31	14.53	25%
1700	20.71	33.52	12.81	62%
1800	5.44	13.03	7.59	140%
December puts				
1400	1.78	5.81	4.03	226%
1500	12.00	22.18	10.18	85%
1600	44.20	58.74	14.54	33%
1700	105.98	118.79	12.81	12%
1800	190.54	198.13	7.59	4%

Figure 6-14 Gold eight-week (40 trading days) historical volatility.