14 Option Pricing Models

You should probably read all the previous chapters before this one.

Option pricing models are mathematical equations which attempt to predict the future price of an option based on inputs available (or guessable) today. The mathematics behind such models is forbidding, but don't worry — I'll be gentle. And you don't really need to understand the mathematics. Most people don't know how computers and television sets work, but are still able to use them. It's the same with mathematical models. What you do need to understand are the uses and limitations of your computer, television, or mathematical model.

In this chapter we will examine the **Black-Scholes model**, which is the best known and most commonly used pricing-model for options. It was first presented by Myron Black and Fischer Scholes in 1973, at the dawn of the options market in the U.S. They received the Nobel prize for this in 1997. There is some dispute about the originality of their work, because a largely similar model was developed in 1908 by the Italian mathematician, Vinzenz Bronzin.

14.1 Volatility and Probability

Like most people, you probably have an intuitive understanding of **volatility**. You understand that it comes and goes, and that when markets are volatile, prices move greater distances and with greater frequency. The wild gyrations of a volatile marketplace may be random, but random motions of a large number of objects can sometimes form common patterns.

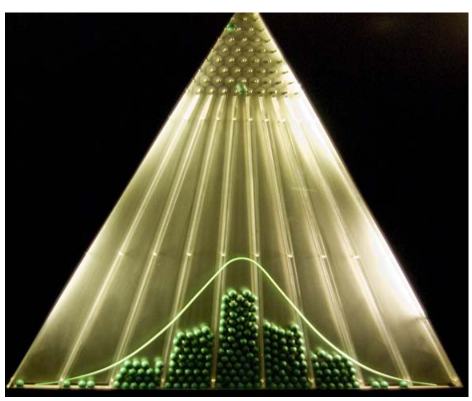

Figure 14.1: A Galton Machine.

Figure 14.1 is a picture of a “Galton Machine”, a famous experiment that every first-year statistics student has seen. It’s a simple machine, and you can make one at home, whiling away many a pleasant hour in solitary contemplation of randomness.

To use it, you drop balls onto the pins at the upper center of the device. When the ball hits a pin, there is an equal probability (50/50) of the ball moving to the left or to the right. While a few balls make it all the way to either side of the device, most tend to accumulate in the middle, and form the classic bell-shaped curve shown in the diagram. The way the balls distribute themselves is called “the normal distribution” by statisticians, and is found frequently in natural phenomena.[1] The shape of the graph isn’t too surprising — most of the balls will end up close to the point they started at. As you move farther towards the left or right of the device, there will be fewer and fewer balls.

By changing the arrangement of the pins, you can alter the “volatility” of the balls. If you increase the volatility, more of the balls will make the distant journey to either end. If you decrease the volatility, more will cluster at the center. You can see this in figure 14.2. For the highest volatility curve (blue line), the “tails” of the curve are much fatter, to accommodate the additional number of balls that have accumulated there. Since more of the ball population has now migrated to the edges, there are fewer at the center of the graph so it must be lower. The situation is the opposite with the low volatility graph — most of the balls have clustered around the center, pushing the peak of the “bell” much higher.

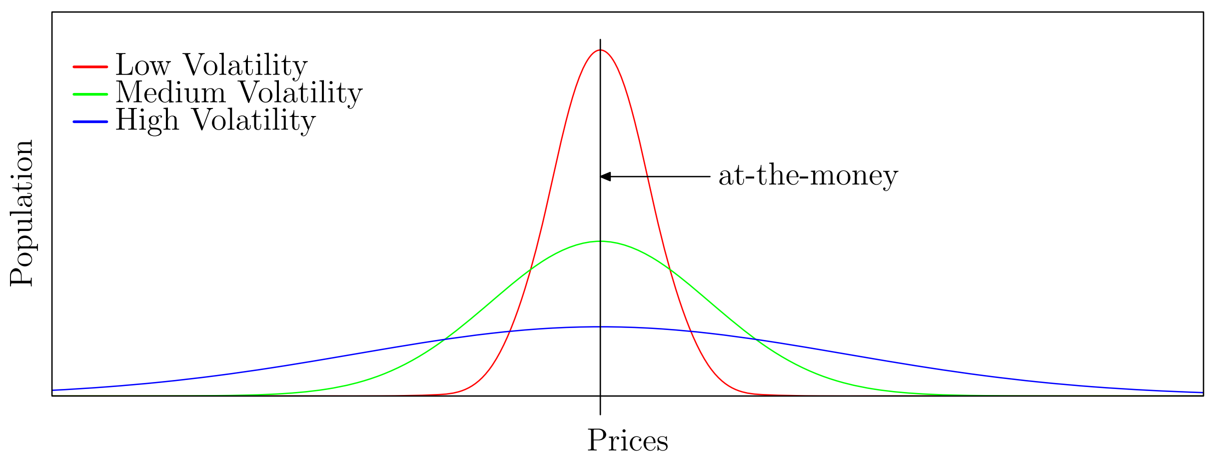

Figure 14.2: Comparing Volatility Distributions.

To convert balls to options, imagine stock options being sprinkled on their underlying, which is right at the center of the bell. They are then buffeted by various market forces, investor sentiment, and who knows what, and eventually settle to the bottom of the graph.

[1]Sir Francis Galton was quite an interesting man. A distant cousin of Charles Darwin, he made significant contributions in many areas. He coined the phrase “nature versus nurture”, did the first research in classifying fingerprints, studied meteorology and was also the founder of the eugenics movement.

You might think that since the graph is so symmetrical, the increase of volatility would have no effect on an option's price. After all, increased volatility makes a large upward movement more likely, but it also makes a large downward move more likely. But there is an important difference between options and their underlying stocks. Only options which are in-the-money will have any value — their loss is limited. So for example, if you are sprinkling calls on the underlying, all of the ones left of the at-the-money line will be worthless at expiration. Only the ones on the right will be worth anything. The value of the stock depends on where the peak of the curve is located, while the value of its options depends on how far the curve spreads out.

14.1.1 The Mean and the Standard Deviation

Any normal distribution is completely defined by two numbers, the **mean** and **the standard deviation**. If you know these two numbers, then you know all you need to know about the distribution.

The **mean** of a distribution will always be the highest point of the curve — it's just another word for 'average', and it shouldn't be surprising that most of the stock option 'balls' end up there. The Black-Scholes model works, in part, by making the stock-price the mean of a distribution.

The **standard deviation** is less familiar, but not difficult to calculate. I won't give the formula here (it's in any statistics text), because it's much more important to understand how it affects you as an investor than it is to learn how to compute it. It is essentially a measure of how "spread out" your population is.

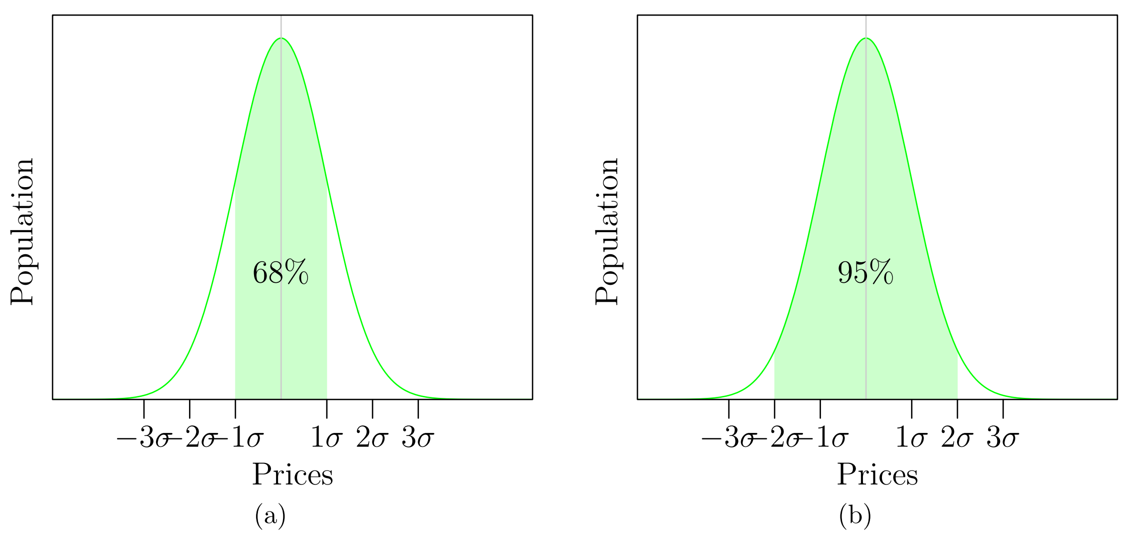

Figure 14.3: (a) 65% of all events occur between the first standard deviations. (b) 95% of all events occur between the second standard deviations.

Figure 14.3(a) and 14.3(b) show the medium volatility curve with the standard deviations marked off, and "probability population" highlighted. The standard deviations are labeled with their usual symbol, σ (pronounced "sigma"). The standard deviation not only tells you how far the distribution will spread out, it also tells you how likely it is something will end up where. The precise number can be calculated, but for all normal distributions:

± 1 Standard deviation contains about 2/3'rds (68.3%) of all the possible occurances.

± 2 Standard deviations contains about 19/20 (95.4%) of all possible occurances.

± 3 Standard deviations contains about 369/370 (99.7%) of all possible occurances.

The mean of this distribution is eight — this just means if you add up all the occurances, and divide by the number of total occurances, you end up with eight. The standard deviation is three.

How much would you have to be paid on a bet that one of your options would end up deep-in-the-money — somewhere past two standard deviations away from the mean? Let's say if that happens, you would make $30 for every $1 you spent. Would it be worth the investment? In this distribution, two standard deviations away from the mean is 6. Since the probability of getting a "hit" within two standard deviations is 19/20, the probability of getting one outside of this area is 1 in 20. But remember that this one chance in twenty includes *both* ends of the curve. So half of the time your option ended up past two standard deviations away it would be *out-of-the-money*. So your option's chance of ending up deep-in-the-money is 1/2 of 1/20, or one chance in forty. Having a pay-off which is less than your risk is not a very good bet to make.

Because the normal distribution is so common, it has been deeply studied for a very long time. As a result, formulas have been developed which make it possible to compute the probabilities associated with every point along a normal distribution curve, as well as the area under various portions of the curve. The Black-Scholes model makes the assumption that the prices of a stock are normally distributed, and uses this distribution to solve for an option's theoretical value, by multiplying each one of a possibly infinite outcomes by its associated probability.

The actual distribution used by the Black-Scholes model is a **lognormal distribution**. It is very similar to the normal distribution, but has fatter tails. One of the principle reasons it was adopted is because the normal distribution allows for the possibility of negative prices — an impossibility in the real world.

14.2 The Black-Scholes Model

Having given you some of the background behind the model, let's take a look at the real thing using call options. At its heart, Black-Scholes calculates the value of a call option as follows:

Stock Price times the probability of the option ending up in-the-money.
Less the present value of the exercise price paid at expiration.

That seems easy enough. I'd give the equation here, but I don't want to frighten you away. I'll save that for the very end. You don't really need to know it. What you do need to know about Black-Scholes is that in order to work, it relies on you to give it the reliable inputs. If you feed it garbage, it will cheerfully give you garbage back.

The original Black-Scholes model needs six inputs in order to predict the price of an option. Here they are:

- The Stock Price (p).
- The Strike Price (s).
- The Risk-free rate (r).
- The Time to Expiration (t).
- The Volatility (σ).

Does this look familiar? It should — these are the same factors we identified in section 1.2 as the most influential on an options premium. The only one missing is the effect of dividends. This is a blemish of the model (and not its only one), but there is an adjustment you can make that helps with this. We'll deal with that later.

The model expects the time to expiration to be entered as a fraction of a year. This is easy and most software will compute the fraction for you given the correct date. The risk-free rate is a little tougher. It should be adjusted to the time-frame of your investment. If you have a 6-month call, the appropriate rate to use is probably the 6-month treasury rate. If you have a 2-year LEAP, then you would use the current rate of the 2-year treasuries. You can find these from free online data services such as Bloomberg.

The really tricky one is the volatility. Options aren't issued with little stickers that say: "Volatility: 30%, guaranteed or your money back." One approach is to use the "historical volatility", which makes the dubious assumption that the volatility of yesterday will be the volatility of today. You can try to "forecast" volatility. This takes the dubious assumptions of historical volatility, and "improves" it by making a guestimate of the future. Good luck with that.

You can also give the Black-Scholes model the current price of an option, the current stock price, the current risk-free rate, and so on, and ask it to tell you what volatility would make all of this possible. This is called **implied volatility**, and is probably the most common approach.[2] This is probably the most common approach, and we'll give an example of how to do this in the very next section.

[2] Some people think this is using the wrong number in the wrong equation to get the right answer. People are very conflicted about Black-Scholes.

14.3 Practical Example: Early Exercise of Put

One thing that is often a concern to investors who have a short put is early assignment. Perhaps the short put is part of a spread that they hope will be profitable, and they aren't really interested in owning shares of the stock. Or perhaps they sold the put for income, and now the put is in-the-money. If you're in a situation like this, you can use Black-Scholes to estimate what the price of the the stock will be when you might need to take action.

The first thing you'll need is a Black-Scholes calculator. If you can run the gnumeric spreadsheet, which is free and really good, this problem is easy to solve. I've provided one here:

http://www.ssr.com/sdb/BAILOUT/Spreadsheets/Black-Scholes.gnumeric

It should also work in Excel. Open it up, and let's get started.

Make your edits in the light-green cells in the input section. You can change the headings (Option 1, Option 2, etc.) to something more descriptive if you want.

In the "Inputs" section, go to the "Puts" sub-section, and make the following entries:

Put Option	Put 1
Data Entry Date	March 20, 2010
Expiration Date	March 4, 2012
Stock Price	10.24
Strike Price	12.50
Risk-free Rate	0.85%
Volatility	

A word about the "Data Entry Date" — this should be the date that the Stock Price and Risk-free Rate are valid on. It defaults to the current date — that will be wrong for this exercise, so just type over it with the date given above. Since this is obviously a 2-year LEAP, we used the current rate paid by 2-year treasuries, which is 0.85%. Make sure you type either 0.0085 or 0.85%, so the spreadsheet doesn't use the wrong value.

Now we need to decide what to do about volatility. Since we don't know what it is, we will hold these inputs steady, give the spreadsheet a 'target' price for the option, and ask it to tell us what volatilty would get us to this set of inputs. To do this, we will need the current price of the option, so looking at your brokers quote screen, you see:

Symbol	bid	ask
Jan-2012 $12.50 Put	$3.10	$3.20

We'll use the average price — $3.15 Now go to the "Tools" menu, and select "Goal Seek". This will pop open a window, where you will make the following entries:

Set Cell:	G21
To Value:	3.15
By Changing:	B26

G21 is the cell address of the Theoretical Price of the option in the OUTPUT section. 3.15 is what we want it to equal, and B26 is the address of the volatility cell.[3] Then click 'Apply'. Magically, a value of 31.40% should appear in the "Volatility" cell you selected. Select "Close" (not "Cancel"), to keep this value for the volatility.

For the curious: Black-Scholes can't actually be solved for the volatility. So Gnumeric (and Excel) use an iterative process to estimate it. First it makes a guess at what it the answer might be, runs the equation with this guess, and compares the result to the theoretical price of $3.15 you entered. If gnumeric's guess at the volatility resulted in a theoretical price higher or lower than $3.15, it modifies its guess, and repeats (iterates), until the guess is "close enough", or until it is apparent that it can't be done.

Now we are part way to our goal. Remember we are trying to find out what price might encourage the owner of our put to exercise on us. Well, what price would that be?

Recall that since this stock pays no dividends, a rational counterparty will only exercise the put when there is no time-value remaining in the put. Why? Because by exercising and having us pay for the stock, they are getting only intrinsic value. They could sell the put for (intrinsic value + time value) and receive more than if they just sold the stock for intrinsic value.

So now we do a *second* "Goal Seek": this time the "Set Cell:" address should be the address of the Intrinsic Value cell. The "To Value:" box should be zero, or perhaps a penny or two, because perhaps your counterparty would sacrifice a penny, And the "By Changing Cell:" box should contain the address of the stock price cell in the "Input" section.

Doing this with a value of "0.01" in the "To Value:" cell gives a stock price of $7.18 as the one which, if reached, may cause your counterparty to exercise early.

14.3.1 Underlying Assumptions

Of course, there is a big assumption underlying all of this; namely, that the inputs for volatility and the risk-free rate remain constant over this extended period. Of course they won't. The risk-free rate will do whatever the risk-free rate does. Should it go up, option prices will increase for both puts and calls. And this has less of an impact on option prices than volatility.

Volatility is the 500 pound gorilla, here. What you generally see is that as a put falls further into-the-money (as long as it is not *extremely* deep), implied volatility will *drop*. By playing around with the various inputs, you can find out that if implied volatility were to fall to 30%, the early exercise stock price will rise to $7.39.

[3] Most spreadsheets allow you to simply click on the cells, and will put the address into the dialogue box. You may find this easier.

On the other hand, if volatility in the market rises, then perhaps the implied volatility would also rise. An implied volatility of 34% today would drop the early exercise stock price to $6.79 today.

And fiddling further with the inputs shows that this price level will change as time passes — it will move *up*. So if it's one year from now, with an unchanged stock price, and the yield curve remains the same as it is today (it won't, but this is just a thought experiment), the risk-free rate will be about 0.34%. And we'll also assume that the volatility stays the same. Going through the steps above with these new assumptions causes the exercise stock price to reach $7.34. By July-2011, it's at $8.00. By October-2011, it's almost to $9.00 and so on.

Play around with it a bit before you move on. Try different values and see how that changes the price of the option. Try some more goal seek scenarios. Have fun.

14.4 Problems with Black-Scholes

Any mathematical model is, by definition, imperfect. Some things are impossible to quantify (such as investor sentiment) or predict (such as macro-economic events). At best, models provide only a partial picture of reality. Many people think Black-Scholes does a poor job, even considering the limitations inherent in all models. Black-Scholes makes many assumptions, all problematical. Here's a list:

1. Options are not exercised prior to expiration.
2. Markets are frictionless.
3. Interest rates are constant over the option's life.
4. Volatility is constant over the option's life.
5. Trading is continuous, with no gaps in the price changes of the underlying.
6. Volatility is independent of the price of the underlying.
7. Over short periods, the percentage price changes in the underlying are normally distributed, resulting in a lognormal distribution of underlying prices at expiration.

These limitations expose users of Black-Scholes to definite risks:

1. It underestimates extreme moves in the market (called "tail-risk").
2. It assumes a frictionless market, giving "liquidity risk".
3. The assumption of a constant volatility gives "volatility-risk".
4. The assumption of continuous prices gives "gap-risk".

Black and Scholes were two of the key players in the Long Term Capital Management disaster.[4] But the magnitude of the disaster, which was a "tail-risk" disaster if there ever was one, sometimes causes us to overlook the fact that Long Term Capital Management also made gobs and gobs of money. So in fact, the mathematical modeling does appear to work, at least some of the time. The danger lies in regarding the model as infallible.

My own approach to using Black-Scholes is adds another layer to my personal approach to options. First I try to understand the business, and layer option trades based on this understanding. On top of this, I'll run some scenarios through a Black-Scholes calculator to see how that looks. If Black-Scholes predicts a favorable outcome, so much the better for me. If Black-Scholes predicts a strongly unfavorable outcome, this might cause me to reconsider my investment. However, I would never make an investment based on a favorable prediction from Black-Scholes that ran counter to my understanding of the business.

Not every investor agrees with this approach, and many use Black-Scholes exclusively, ignoring the fundamental business. They seem to make money doing this, but I am not comfortable with it. Whatever your personal approach, running various scenarios through Black-Scholes can be very educational, particularly for investors new to options.

14.5 Extending the Model: Incorporating Dividends

One strange assumption that Black-Scholes makes is that the underlying does not pay a dividend. As has been discussed previously, dividends tend to lower the price of call options, and increase the price of put options. While the effects on short-term options are slight, the effects on long-term options such as LEAPS can be substantial, possibly substantial enough to turn a wining position into a losing one.

Fortunately this blemish is easy to fix. There are two methods:

1. Deduct the present value of all dividends that will be paid during the life of the option from the stock price input to the Black-Scholes model.

2. Assume the company pays a consistent dividend yield. Divide the stock price by the exponential function of this yield, multiplied by the time to expiration. The resulting mess is then used instead of the stock price.

Neither method is perfect. The first suffers from future dividends being unknowable — what if they raise or lower them? The second method suffers from over simplification. In a world where an average stock price can fluctuate in 30-50% in value during a single year, a constant yield assumption strikes me as bizarre. As you may have guessed, I am an advocate for the first method, and that is the one that is implemented in the Black-Scholes calculator I distribute.

Start your handy Black-Scholes calculator, and make the entries shown in figure 14.4. Note that we've also changed the column headings in this example — just to remind us

[4] Great book on this: "When Genius Failed" by Roger Lowenstein

Call Option	With Dividend	Without Dividend
Data Entry Date	November 20, 2009	November 20, 2009
Expiration Date	January 20, 2012	January 20, 2012
Stock Price	28.00	28.00
Strike Price	20.00	20.00
Risk-free Rate	1.05%	1.05%
Volatility	32.5%	32.5%
Quarterly Dividends Per Share		
1	$0.13	
2	$0.13	
3	$0.13	
4	$0.13	
5	$0.13	
6	$0.13	
7	$0.13	
8	$0.13	
Total Dividends Received	$1.04	0

Figure 14.4: The effect of dividends on options.

what is going on.

Once you've done this, you'll notice a radical change in the theoretical prices of the options predicted by the model. The option with dividends will have a theoretical price of $9.04. The same option without any dividends will have a theoretical price of $9.89. As the stock price increases, the difference becomes more extreme. For example, at a share price of $100, the model predicts a theoretical price of $79.42 with dividends, and $80.45 without. That's a difference of over $100 per contract!

Still, this is not a loss. But consider what can happen if you were to enter a diagonal spread, selling short-term calls against your long term leap. If you had sold a $27 call when the shares were down, and they moved to $100 at the short calls expiration, your short call would be worth $73.00 (Stock price - strike price of short call = Intrinsic Value). But since you sold it short, that's money you *owe* — you would need to buy the call back. So your profits now look like this:

With Dividends $79.42 - $73.00 = $6.42.

Without Dividends $80.45 - $73.00 = $7.45

Add in commissions, and the cost of the LEAP (which was probably around $7.00), and you very well might be looking at a loss in this situation.

Running such "what-if" scenarios is an excellent application for the Black-Scholes model.

14.6 Grand Finale: The Black Scholes Formula

Now that we have reached the end, as promised, here is the Black-Scholes model, in all its glory.

$$C = pN(d_1) - se^{-rt}N(d_2)$$
$$P = -pN(-d_1) + se^{-rt}N(-d_2)$$

where

$$d_1 = \frac{\ln(p/s) + (r + \sigma^2/2)t}{\sigma\sqrt{t}}$$
$$d_2 = d_1 - \sigma\sqrt{t}$$

- C = theoretical value of a call.
- P = theoretical value of a put.
- p = Stock price.

- $N(x)$ = cumulative normal density function.
- s = strike price of option.
- e = Euler's number (base of the natural logarithm).
- r = risk-free interest rate.
- t = time to expiration in years.
- σ = volatility.

For those who have have journeyed this far with me, thanks for the company. Please don't use what you have learned in this book to make all the money in the world — leave some for the rest of us.

And thanks for listening.

14.7 Chapter Glossary

Standard Deviation Loosely speaking, this is a mathematical description of the average distance from the average of a distribution. It tells you how spread out your population is.