

GraefCrystal.com

The Crystal Report on Executive Compensation



The Black-Scholes Model Seems to Overvalue Options

by Graef Crystal

February 8, 2010

Warning: This article won't be of interest except to people who care deeply about the subject of option valuation.

Now that you have been warned, let's start the discussion with the so-called Minimum Option Value Model (MOVVM).

MOVVM

The formula for determining an option's present value using the MOVVM is comparatively simple. The present value equals the Fair Market Value at Grant less the Present Value of the Strike Price. And if the expected future dividend is greater than zero, an additional term is added, which adjusts the Fair Market Value at Grant by subtracting the Present Value of the Expected Yield.

To simplify matters, let's assume that the expected yield is zero. Then let's make the following additional assumptions:

- The market price at grant is \$50.00.
- The strike price is the same \$50.00.
- The term of the grant is seven years.

- The risk-free rate for a zero-coupon Treasury bond (a so-called “strip”) with a seven-year maturity is 4.2 percent.

On the basis of the above, the present value of the option is: $\$50.00 - \$50.00 / (1 + \text{risk free rate of } 4.2 \text{ percent, or } 1.042)$ to the power of the option’s term, or 7. The result is \$12.51.

Now let’s set up two hypothetical investors, each investing \$50.00. Investor A buys one share of stock, paying \$50.00. Investor B buys one option worth \$12.51 and then invests the balance of his \$50.00, or \$37.49, in a Treasury strip carrying interest of 4.20 percent.

Start the clock running. At the end of the seven-year term of the option, there are three possible outcomes:

- The stock will be selling for more than \$50.00.
- The stock will be selling for \$50.00, the same price at which we started seven years earlier.
- The stock will be selling for less than \$50.00.

If the stock is selling for more than \$50.00, both Investor A and Investor B end up the same, no matter how high the future price.

For an illustration, let’s assume that at the end of the option’s seven year term, the stock is selling for \$200 a share. Investor A ends up holding one share of stock worth \$200. Investor B cashes in his Treasury strip. At an interest rate of 4.2 percent, his bond, for which he initially paid \$37.49, is now worth \$50.00. He thereupon takes his \$50 and uses it to exercise his option. At the end of this transaction, he, too, is holding one share of stock worth \$200.

Now let’s assume the stock doesn’t move, and the option’s term ends with the stock selling for \$50. Investor A has \$50.00 from his one share of stock. Investor B’s option expires worthless. But his \$37.49 investment in a Treasury strip is now worth \$50.00. So both investors come out the same.

Finally, let’s assume the stock drops to \$25 during its seven-year term. Investor A ends up with one share worth \$25.00, for a loss of \$25.00. Investor B’s option is worthless, but his Treasury strip is worth \$50.00. So he comes out better than Investor A.

If that’s the case, then why shouldn’t Investor A go the same way as Investor B and buy one option and invest the remainder of his \$50.00 in a Treasury strip?

Enter the Black-Scholes

Before answering this question, consider that the Black-Scholes option pricing model uses six assumptions to produce a present value:

- The market price at grant.
- The strike price.
- The term of the grant.
- The risk-free rate.
- The expected dividend yield.
- The expected volatility.

The MOVVM uses five of these six assumptions – all except the last one, the stock's expected volatility.

Or to put it more accurately, the MOVVM uses all six assumptions – the first five plus a hidden assumption, namely, that the expected volatility will be 0 percent.

But the probability that the actual future volatility of a stock will be 0 percent is infinitesimally small. A volatility of 10 percent is about as low as you'll ever encounter.

The introduction of a real-world expected volatility assumption causes the option's present value under the Black-Scholes model to be greater than the figure produced by the MOVVM.

As a result, Investor B will have to pay more for his option than \$12.51. And because he must do so, the amount he can invest in a Treasury strip will be comparably smaller. Therefore, if the option finishes underwater, he, too will pay a penalty, because his Treasury strip will not be worth \$50.00.

An Example

Turning to the real world, let's put those assumptions to work:

- Market price at grant of \$50.00.
- Strike price of \$50.00.
- Term of seven years.
- Risk-free rate of 4.2 percent.
- Expected dividend yield (for the sake of simplicity) of 0 percent.

- And a real-world expected volatility assumption of 40 percent.

According to the Black-Scholes model, that package of assumptions will produce a present value of \$24.49 a share.

Now let's open the hood of the Black-Scholes model. One of the key things powering the model is a normal distribution of likely stock prices at the end of the option's term. Or, more precisely, a lognormal distribution. Stock prices, like pay itself, produce a higher average than median. But a normal distribution produces an average that is the same as the median. To correct for this skewness, logarithms are used, because when the logarithms are displayed on a chart, the distribution looks to be quite normal.

Now let's construct 10,000 lognormal scenarios for the future price of the stock. Taken together, these 10,000 scenarios produce a perfectly normal distribution of stock prices at the end of Year 7.

But underlying that distribution of 10,000 scenarios is yet another assumption.

Using the so-called Capital Asset Pricing Model (CAPM), the expected total return on the stock will be:

- The risk-free rate of 4.2 percent.
- A so-called equity risk premium. Using recent Bloomberg data, the equity risk premium I postulated was 6.2 percent.
- The Beta on the stock. I am assuming that the beta is 1.00, meaning that the stock will rise and fall at the same percentage as the overall market.
- The formula for the CAPM is Risk Free Rate + Equity Risk Premium X Beta, or in this case, 10.2 percent.

Given all those assumptions, the average ending stock price for all 10,000 scenarios will be \$99.35.

So how do I know if I'm right about this? Well, some years ago, my mathematically-gifted erstwhile Towers Perrin colleague, Steve O'Byrne (who now has his own consulting practice, called Shareholder Value Advisors) ran into Myron Scholes at a conference, and the Great Man gave him a formula for extracting the expected future value of an option from its present value.

After correcting for the fact that the option is worthless for any ending price below \$50.00 and for the fact that an optionee will have to pay \$50.00 to exercise his option under those scenarios where the ending price is higher than \$50.00, my math produces an expected future option value of \$60.11 for all 10,000 lognormally-distributed scenarios. And Professor Scholes value turned out to be \$60.09.

A Comparison

Now, let's go back to Investor A and Investor B.

Investor A buys one share of stock, paying \$50.00. Averaging all 10,000 scenarios, his one share of stock at the end of Year 7 will likely be \$99.35.

If you examine those 10,000 scenarios, you will find a wild range of possible outcomes – from a mere \$0.62 a share at the end of Year 7 to a value of \$5,178 a share.

That variation incorporates a standard deviation range of negative-4.265 deviations to positive 4.265 deviations. The likelihood of a stock deviating by more than that figure is only around 1 in 100,000.

So let's truncate the range to negative-3.00 standard deviations to positive-3.00 deviations. That still covers more than 99 percent of all likely outcomes.

On that basis, the range of possible outcomes shrinks to \$2.41 a share to \$1,345 a share, to my mind a more real-world range. And the average expected outcome drops to \$97.02 from \$99.35. The number of scenarios in my model decreases to 9,973.

So what about Investor B? He buys one option on the stock, paying \$24.49 a share. Then he invests the remainder of his \$50.00, or \$25.51 in a Treasury strip bond paying an interest rate of 4.2 percent a year. We will assume here that the proceeds of the option, if any, can be attained through a cashless exercise.

At the end of year 7, that Treasury strip bond will be worth \$34.03. So if the stock finishes at \$50.00 or less, Investor B will lose money. Indeed, he will lose money until the stock has risen to \$65.97. At that point, Investor B has \$34.03 from his Treasury strip and \$15.97 of appreciation in his option, for a total of \$50. That is his breakeven point.

If we average all 9,973 possible outcomes, Investor B ends up with an average value of \$91.77.

But that result of \$91.77 is substantially lower than the average outcome for Investor A, which was \$97.02.

So the way I see it, Investor B should forget about his fancy Treasury strip and option combo and just pay \$50.00 and buy a share of stock.

So how can we equalize the two outcomes, so that Investor B also has an average value of \$97.02 a share?

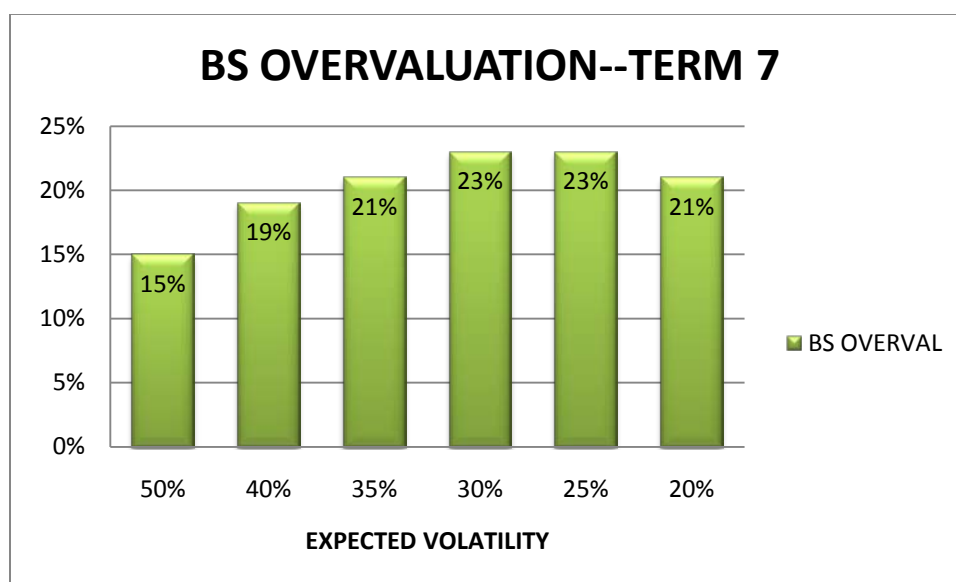
We can achieve that result if Investor B pays less than \$24.49 for his option, thereby giving him more to invest in a Treasury strip.

We can equalize the outcomes for Investors A and B if Investor B pays \$20.56 for his option, a figure that is 19 percent lower than the Black-Scholes value of \$24.49.

Simulations


To provide some further insight, I ran a number of additional simulations, using the same Black-Scholes assumptions as outlined earlier, except for term and volatility.

The following chart shows the percentage by which the Black-Scholes seems to overvalue options as a function of expected volatility. All scenarios involve a term of seven years.



The volatility range chosen – 20 percent to 50 percent covers the vast bulk of possibilities. An expected volatility of 20 percent is equal to the 2.7th percentile of implied one-year call volatilities for companies comprising the Standard & Poor's 500 Index. An expected volatility of 50 percent weighs in at the 93rd percentile.

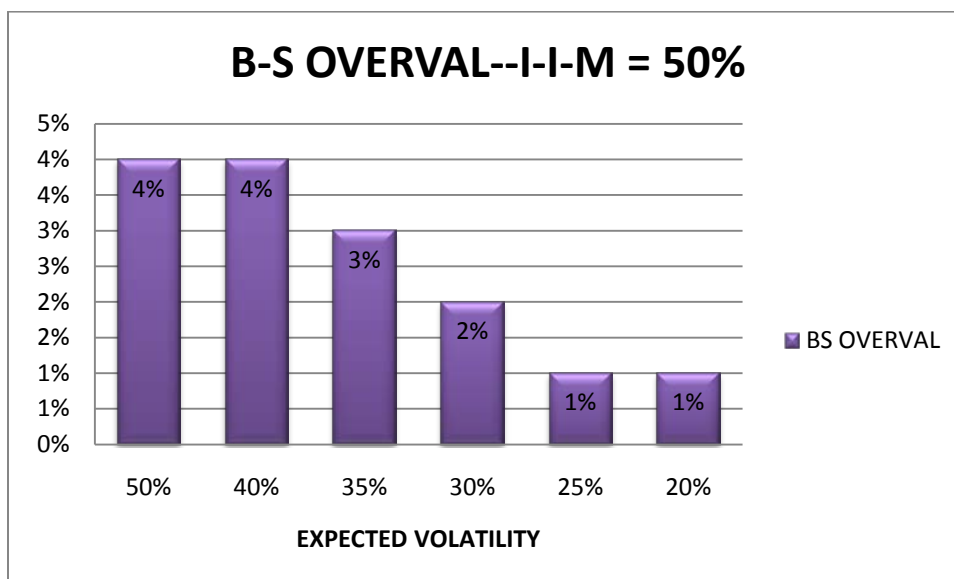
This next chart shows the extent of overvaluation, but this time using a 10-year term.



Both of the above charts assume, among other things, that the strike price is equal to the market price at grant.


But what if it is not?

I reverted to a seven-year term and ran another set of simulations assuming that the strike price would be only 50 percent of the market price at grant – in other words, that the option would be “in the money” from the get-go.



This chart shows that the Black-Scholes seems to be fundamentally on the mark if the option is in-the-money sufficiently at the time of its grant.

But just the opposite conclusion emerges when we assume that the option is out-of-the-money at its grant – in this case, that the strike price is 150 percent of the market price at grant.



The degree of apparent overvaluation rises significantly when the option is out-of-the-money at its grant. (I did not include a scenario for a 20 percent expected volatility level, because the result would have gone right off the chart. It makes almost no sense to grant a significantly out-of-the-money option in a company where the expected variance in the stock price is so small.)

Conclusion: I think a credible case can be made for arguing that the Black-Scholes model significantly overvalues stock options – especially those where the strike price is equal to the market price at grant (by far the most common alternative) or higher than the market price at grant.

Steve O’Byrne may or may not agree with the logic in this article. Still, both of us agree that the Black-Scholes overvalues options – even if his rationale may not be the same as mine. He is fond of pointing out that underpinning the Black-Scholes is an assumption that the investor is perfectly diversified. That would be considered, to use a favorite phrase of Charles O’Reilly, “an heroic assumption”. (Professor O’Reilly was a former colleague when I was teaching at U.C. Berkeley. He is now a professor at the Stanford Business School.) For the typical executive, the largest assets relative to his net worth are the equity in his house(s) and the stock he owns in the company. That hardly makes him perfectly diversified. Indeed, if I recall correctly, the discount that Steve thinks should be applied in an extreme case of non-diversification is somewhat north of 50 percent.

Graef Crystal's is now in his 51st year in the executive compensation field. He has been a director of compensation for General Dynamics and Pfizer, worked as a consultant for Booz, Allen & Hamilton, served as worldwide practice director at Towers Perrin for 18 years, was a professor at the University of California at Berkeley's Haas School of Business for 10 years and a syndicated columnist for Bloomberg News for almost nine years. (He still contributes occasional columns to Bloomberg News.) He has written six books and more than 1,600 articles on executive pay.