WHITE PAPER
MAY 2008

The Cost of Black-Scholes

The Assumptions Underlying the Famous Options-Pricing Model Are Notoriously Unreliable. Here's How Their Failure Could Affect an Options-Replication Strategy

Nick Mocciolo
FSA, FRM, Senior Vice President

Asset managers often try to put a floor under the price of stocks they hold, typically by buying "put options," which entitle the holder to sell a stock at a certain price by (or on) a certain date. But there is another way to hedge portfolios that can, under the right circumstances, provide the same level of protection at a lower cost, albeit with greater risk. That way has come to be known as "portfolio insurance" — the art of continually rebalancing portfolio assets according to moves in the market, exchanging stocks for cash when stocks fall; doing the reverse when stocks rise. The technique is also called "delta hedging" or "dynamic replication."

Contrary to popular belief, dynamic replication did not die with the 1987 stock-market crash, whose skids it famously greased. It remains a tool of asset managers seeking to protect portfolio values. But how does one decide whether to buy options or to replicate? Buying a put option entails a rather straightforward risk/reward tradeoff: You pay a premium and receive the stipulated protection. It's a "set and forget" strategy with a fixed cost and known payoff.

By contrast, dynamic replication requires frequent rebalancing and entails unknown costs. Though by definition it does not involve purchasing put options, an understanding of option pricing is crucial to the strategy. An option's price comprises many parts — the price of the underlying stock, the option's strike price, time to expiration and others — so valuing the option requires some sort of pricing model. Under the most widely used model, known as Black-Scholes,[1] the price of the put option is the same as the theoretical cost of implementing a dynamic-replication strategy that would afford the same payoff as the option. In other words, under Black-Scholes assumptions, these two strategies are interchangeable (see *Option Pricing via Dynamic Replication* on pages 9 and 10). In fact, it was while examining a replication-based hedging strategy that Black and Scholes discovered their now-famed formula.

The problem with Black-Scholes, as many have observed, is that it makes several critical assumptions which in real-world financial markets don't always hold up. Indeed, they never hold up as a set. One of those assumptions is that the underlying stock's volatility — roughly, the degree of fluctuation in its price — is constant and known with certainty. In reality, volatility is the one unobservable component of an option's price, so options traders look instead at "implied volatility," the unique volatility level for the stock suggested by the price of the option (as explained on pages 9 and 10, there is a one-to-one relationship). Since actual replication costs are heavily influenced by the volatility of the underlying stock, the level of volatility implied by the price of a given option is an important element in assessing the likelihood that a replication strategy will outperform an options-based strategy. Thus, a critical consideration in the buy-vs.-replicate decision is the relationship between an asset manager's expectation of the stock's volatility over the timeframe covered by the option and the implied volatility extracted from the option price.

[1] Economists Fischer Black (1938-95) and Myron Scholes (1941-) introduced the model in a 1973 article titled "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81 (3): 637-654.

Contrary to popular belief, dynamic replication did not die with the 1987 stock-market crash, whose skids it famously greased.

But there are additional considerations, such as the reliability of the other assumptions embedded in the option pricing model. While it's old news that replication entails risk, it's not widely understood how severely each Black-Scholes assumption can impact the effectiveness of the strategy. As you might expect, the more that markets depart from Black-Scholes, the greater the potential divergence between the price of a put option and the price of replicating it — prices that in a Black-Scholes universe would be identical. That divergence can be a source of savings, but also of significant risk. Here, we try to quantify that risk through simulations that relax key Black-Scholes assumptions – first singly, then ultimately as a group – and measure the effect on a hypothetical replication strategy.

It's probably impossible to make unconditional statements about the frequency and magnitude with which replication will outperform put buying, owing to the multitude of factors that ultimately influence the outcome. But our analysis does demonstrate some important themes. In general, at least under the conditions of our hypothetical replication experiment, it shows that replication is likely to be cheaper. In our experiment, this was the case 78% of the time, with an average savings of about 14% vs. the corresponding options-based strategy. This makes sense, as option writers require a premium in exchange for accepting the risk of selling and hedging options, and option buyers may be willing to forgo potential savings to avoid replication risk. However, there is significant variability in savings. In our experiment, the standard deviation of outperformance is about 25% the cost of buying options. There is also significant tail risk: The worst-case replication outcome cost 2.5 times what the options-buying strategy cost.

Critical Assumptions

The Black-Scholes formula makes many assumptions, but for our purposes these are the critical ones:

1. There is no limit to the number of times a replicating portfolio can be rebalanced; in other words, continuous trading is possible.
2. The stock price obeys geometric Brownian motion (meaning that the stock's return over a given period is normally distributed and is not influenced by returns in prior periods).[2]
3. The volatility of the stock price is constant.
4. The volatility of the stock price is known with certainty.
5. Fractional shares of stock can be traded.
6. All investors can borrow and lend at the risk-free rate.
7. The risk-free rate is constant.
8. There are no transaction costs when rebalancing the portfolio.

[2] A stock price follows geometric Brownian motion if it can be written as $dS = \mu S dt + \sigma S dW$, where μ and σ are the drift and volatility, respectively, of the stock price. Under this framework the (random) return of the stock price has a normal distribution and the stock price itself has a lognormal distribution.

Figure 1

Interval Between Portfolio Rebalancing	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
None	0	0
15 Minutes	(0)	15
30 Minutes	0	22
1 Hour	(0)	30
2 Hours	(0)	43
4 Hours	(1)	58
1 Day	(2)	82

Figure 2 Simulated Dynamic Replication vs. Put Payoff ($millions)

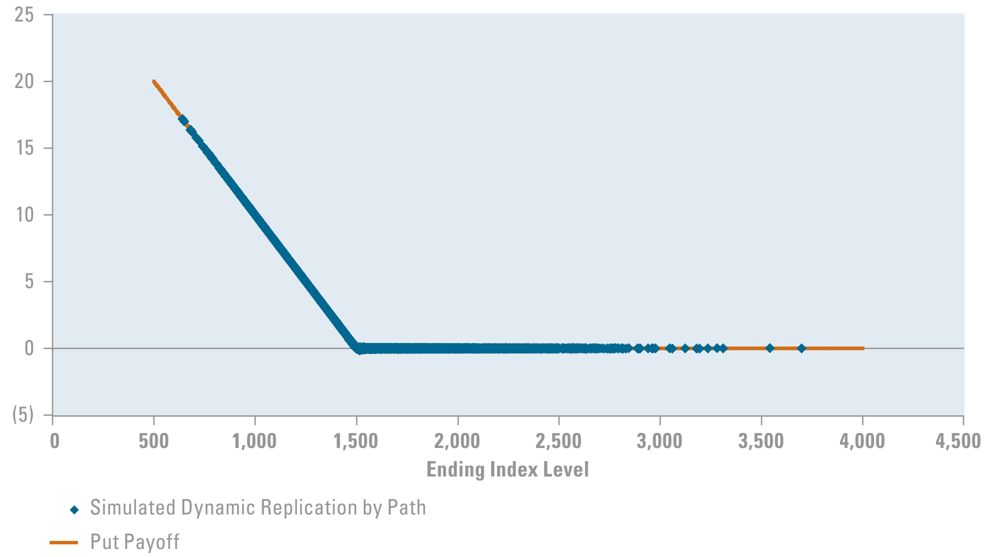

We will try to examine these assumptions by isolating their impact in a simulation-based hedging experiment which assumes that a fund manager wants to protect a portfolio against a potential drop in a stock index. Our hypothetical trade has the following characteristics:

- The initial value of the index is 1500.
- The index pays no dividends.
- The time horizon considered is one year.
- The risk-free interest rate is 5% with continuous compounding.
- The fund manager wants to hedge the portfolio by using 20,000 put options on the index with a strike price of 1500. The options are therefore "at the money" and the fund manager has protection on $30 million of "notional value," the market value of the stock controlled by the option.
- The implied volatility for one-year put options on the index is 25%.
- The actual volatility of the stock over the ensuing one year is 25% (though in the real world we can't know this in advance).

Assume for a moment that Black-Scholes assumptions prevail. Applying Black-Scholes to the parameters above, the value of a single put option would be $111.883925, making the total premium 20,000 x $111.883925 = $2,237,678.50. Alternatively, the fund manager can replicate this option. As explained on pages 9 and 10, we know that under Black-Scholes assumptions these two strategies should produce identical results.

But what if one or more of the Black-Scholes assumptions fails? To try and quantify the impact, we conduct a simulation that samples 5,000 random paths for the index over the one-year horizon, and for each path we calculate the performance of the dynamic-replication strategy described above. Specifically, we assume the fund manager commits $2,237,678.50 to fund the replicating portfolio, and injects no additional cash over the one-year period. In the end, we can compare the replication strategy's payoff with what would have been earned by buying options.

Assumption 1: Relaxing the assumption of continuous trading

Continuous rebalancing, of course, is impractical. But we can get pretty close, as Figure 1 shows. The row corresponding to "15 minutes," the highest frequency tested, puts us in an "almost Black-Scholes" universe: The average outperformance is zero, with a standard deviation of only $15,000 on a notional value of $30 million. In other words, the replication payoff is more or less identical to the options payoff, as you would expect. Figure 2 depicts the hedging performance by path. Each blue dot is the result of one of the 5,000 scenarios, showing the ending value of the replicating portfolio plotted against the payoff of the option being replicated (the orange line). The close alignment of blue and orange confirms the virtually identical payoffs.

Figure 3 Distribution of Simulated Jump Sizes (Probability)

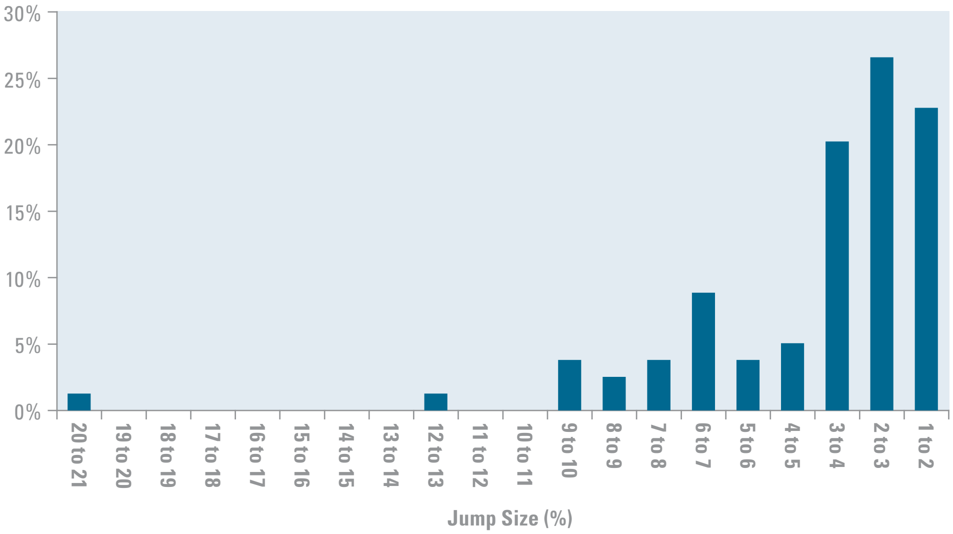

Figure 4

Stock-Price Jumps Implemented	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
No	(2)	82
Yes	(2)	112

Figure 5

Volatility of Stock Price	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
25% Constant	(2)	82
20% / 30% Blend	(0)	117

Clearly, however, as we lengthen the time between rebalancing — that is, the further we get from the Black-Scholes ideal of continuous trading — the less precisely the replication outcome tracks the option outcome. The impact, however, seems limited to the variability of outperformance. That's because a mishedged position can either help or hurt, depending on the index's next move. Over the life of the strategy, at least in our example, good moves tend to cancel bad ones, and vice-versa.

In short, then, violation of the "continuous trading" assumption has a negligible impact on expected outperformance under our hypothetical replication experiment, but introduces some risk to the strategy. (In all subsequent sections, this analysis assumes daily rebalancing, in order to allow fair comparisons and ease computations.)

Assumption 2: Relaxing the assumption of geometric Brownian motion

There are plenty of stock-price processes that differ from geometric Brownian motion. Sticking with simple deviations, we implement a geometric Brownian motion process with one twist: During each randomly sampled path, we introduce one discontinuous jump downward in the stock price on a randomly selected day. We choose the size of the jump from a distribution (Figure 3) created by considering the largest single-day fall in the S&P 500 index in each calendar year from 1928 to 2006 (a total of 79 possible values, each of which might be chosen any number of times). We adjust each overall path to ensure that the total volatility of the index remains at 25%, including the jump.

Figure 4 shows the results. As expected, there is a significant increase in variability of results, but the mean does not change materially. While this may seem counterintuitive, recall that we implemented a jump while holding the overall level of volatility unchanged.

Assumption 3: Relaxing the assumption of constant volatility

We want to relax assumption No. 3 without relaxing assumption No. 4, so we implement a simple procedure to test the impact of variable volatility. We assume that the stock price's volatility is 20% for the first six months and 30% thereafter, constraining volatility over the whole path to 25%. In essence, we assume that the fund manager correctly predicts the total realized volatility over the year, but does not predict it correctly for any period within the year.

Figure 5 shows the results. While the average level of outperformance does not change materially in this analysis, the variability of outperformance changes significantly.

Clearly this is a highly discretionary method of incorporating non-constant volatility into the model. There are many other ways to do this, including introduction of stochastic volatility models. We've chosen this method for simplicity, but, as in the previous section, other methods could produce different results.

Figure 6 Implied vs. Realized Volatility: S&P 500

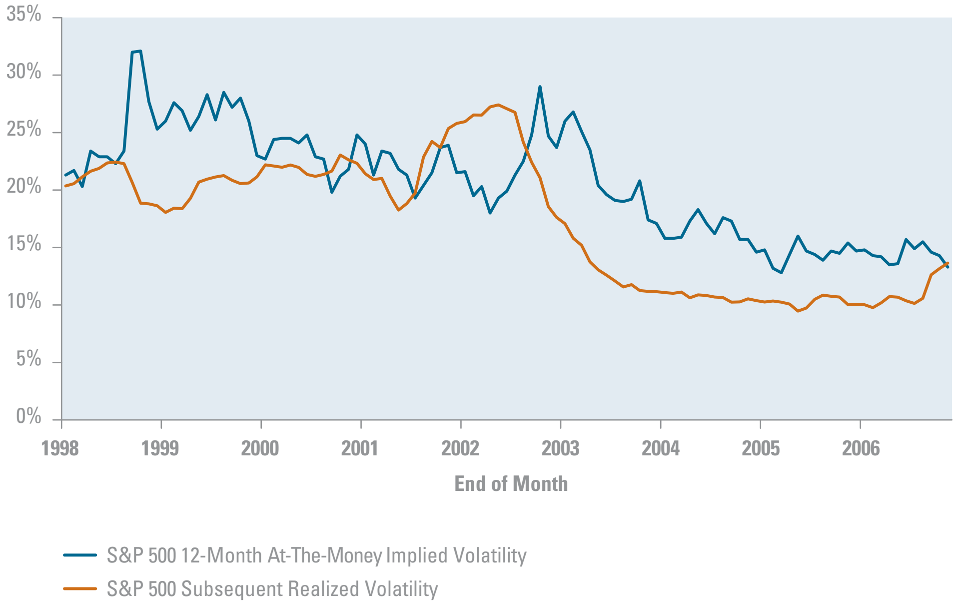

Source: S&P implied-volatility data from Deutsche Bank. Index data from Bloomberg. Realized volatility calculated by Hartford Investment Management.

Figure 7

Realized Volatility Known in Advance	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
Yes	(2)	82
No	399	563

Figure 8

Trading of Fractional Shares Permitted?	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
Yes	(2)	82
No	(2)	83

Assumption 4: Relaxing the assumption of known volatility

It is, of course, impossible to know in advance what volatility a stock will exhibit over an option's life. Implied volatility tends to be a reasonable indicator of future volatility, but there are times when it's a very poor indicator.

Figure 6 shows the historical relationship between implied and realized volatility. The blue line shows the S&P 500's implied volatility for 12-month, at-the-money options for month-end days from December 1997 to October 2006 (107 data points). The orange line shows the actual volatility exhibited by the S&P 500 Index over the 252 trading days following each month-end day.

There are many ways to explore the impact on overall performance when realized volatility differs from implied volatility. As a simple test, we conduct our simulated replication with the following tweak. The 107 data points described above lead to 107 differences between implied volatility and realized volatility over a one-year horizon. On each path, we sample from this distribution and alter realized volatility according to the outcome. For example, for a given path, if we sample a difference of 100 basis points, we then run that path assuming realized volatility will turn out to be 25%−1%=24%. Note, however, that our fund manager still performs replication under an assumption of 25% volatility.

For most paths in our experiment, realized volatility is lower than implied volatility. In many cases this allows dynamic replication to outperform the put option. After all, the pricing of the option assumes some level of rebalancing costs. When realized volatility is low, however, the actual cost of rebalancing may prove lower than what was projected. Higher volatility can have the opposite effect.

The results are displayed in Figure 7. The average outperformance is significantly positive, mainly because — at least for the strike price, option tenor and time period analyzed — implied volatility tends to exceed realized volatility. But notice also the significant increase in variability of results. This is related to the wide variation of realized volatility levels, and is clearly a dominant effect in overall performance, as would be expected from a theory that focuses on replication, the cost of which relies heavily upon market volatility.

Assumption 5: Relaxing the assumption that fractional shares can be traded

Rather than assume that fractional shares are available, we instead posit that the hedging vehicle is a futures contract with a multiplier of 50 on the index. So in essence the theoretically pure mix of cash and stock suggested by the Black-Scholes assumptions at any given time is modified to be consistent with the closest possible match attainable using whole lots of 50 shares. The impact on our experiment is minimal with respect to both the average level of outperformance and the variability of outperformance. Figure 8 displays the results.

The impact of relaxing this assumption on the results of any dynamic trading strategy is likely to be much more material for smaller trades. In general, for institutions, this is likely an immaterial issue.

Figure 9

Borrowing and Lending Spread	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
0bps	(2)	82
25bps	(37)	84

Figure 10

Interest Rate Volatility Assumption	Interest Rate / Equity Correlation	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
0	N/A	(2)	82
100bps	+20%	(15)	121
100bps	0%	(2)	123
100bps	-20%	10	122

Figure 11

Bid / Offer Spread on Shares	Flat Commissions per Trade	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
0.00	0.00	(2)	82
0.50	3.00	(39)	83

Assumption 6: Relaxing the assumption that investors can borrow and lend at the risk-free rate

We assume that investors can borrow at the risk-free rate plus 25 basis points and can lend at the risk-free rate minus 25 basis points. In the dynamic replication of put options, the investor must sell the underlying stock short and invest the proceeds in a bank account. A lending rate that is 25 basis points below the risk-free rate will reduce earnings in the bank account and make the dynamic hedging strategy more likely to underperform.

Figure 9 displays the results. Though there is significant underperformance on average, the variability of results does not materially change. This is intuitive, as the hedging process has the same level of effectiveness in dampening the impact of the index movements, but the average performance is dragged down by the increased funding cost resulting from the reduced bank-account earned rate.

Assumption 7: Relaxing the assumption that the risk-free rate is constant

Here, we introduce a stochastic process for the risk-free rate. It is modeled using a one-factor Hull-White model[3] with an absolute volatility of 100 basis points per annum. Since we are now assuming two random variables — the stock price and the risk-free rate — then we must also consider the correlation between them.

Figure 10 shows the simulated replication results after introducing randomness to the risk-free rate, making various assumptions as to interest-rate/equity correlation. In all cases, the variability of outperformance increases significantly. In addition, the correlation assumption seems to influence the average level of outperformance: The higher the correlation, the worse the average performance.

To see why, consider the following. To replicate a put option on an index that's falling you'd have to short-sell an increasingly large number of shares, in effect adding cash to the replicating portfolio. If rates and stocks are positively correlated, a falling index is more likely to mean falling interest rates. In other words, the bigger the bank account, the lower the interest it'll earn on average, and vice-versa.

This says something important about dynamic-replication strategies for put options: The overall process is likely to be less effective during times of crisis, when a "flight to quality" generates large-scale selling of stocks and buying of Treasurys, forcing stocks and rates down simultaneously. Paradoxically, these are often the scenarios that put-option buyers are most interested in protecting against.

Assumption 8: Relaxing the assumption that there are no transaction costs

Adding transaction costs to the replication process is straightforward, and the results are intuitive. We assume a $0.50 bid/offer spread on all share sales or purchases and a flat commission of $3. We make the same number of trades, but now they're more expensive. This does not change the variability of outperformance very much, but it significantly worsens average performance, as shown in Figure 11.

[3] The Hull-White model assumes that the process for the short rate of interest is $dr = (\Theta(t) - \alpha r)dt + \beta \cdot dW$, where ($\Theta$(t) is the time-dependent reversion target chosen to make the short-rate model arbitrage-free with respect to the initial term structure and α and β are model parameters that govern the strength of mean reversion and volatility, respectively.

Figure 12

Assumption Set Relaxed Simultaneously	Average Level of Outperformance ($000s)	Standard Deviation of Outperformance ($000s)
No	(2)	82
Yes	327	571

Figure 13 Simulated Dynamic Replication vs. Put Payoff ($millions)

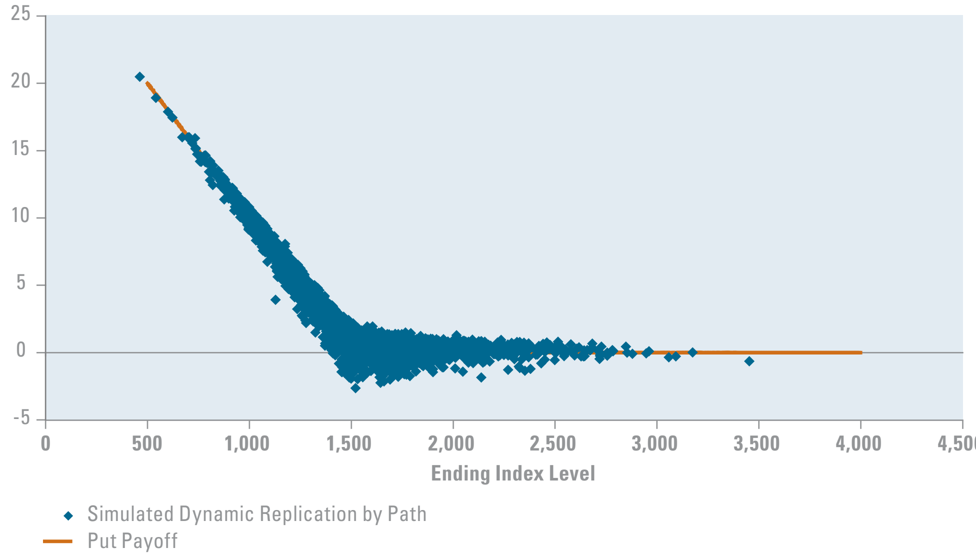

Relaxing Everything

We now put this all together by re-running our simulation relaxing the various assumptions simultaneously. This will give us some insight into the actual buy-vs.-replicate decision that fund managers face when adopting option-like strategies. In each simulation, we do the following:

- Sample to obtain a value for the realized volatility to be exhibited by the index price in the path (relaxing the known-volatility assumption)
- Simulate a random path for the risk-free rate of interest (relaxing the risk-free rate assumption), assuming interest-rate volatility of 100 basis points a year and an equity/interest-rate correlation of 0%
- Simulate a random path for the index price, including a simulated jump and non-constant levels of interim volatility, in a way that achieves the overall level of volatility specified above (relaxing the geometric Brownian motion and constant-volatility assumptions)
- Add 50-share lots only for rebalancing (relaxing the fractional-shares assumption)
- Add a $0.50 bid/offer spread on rebalancing and a flat $3 commission per rebalancing trade (relaxing the assumption of no transaction costs)
- Change the lending rate to be 25 basis points below the (random) risk-free rate and the borrowing rate to be 25 basis points above the (random) risk-free rate (relaxing the assumption that borrowing and lending occur at the risk-free rate)
- Assume daily rebalancing (relaxing the assumption of continuous trading)

Figures 12 and 13 display the results. While the average outperformance is materially positive, owing to the tendency of realized volatility to be lower than implied volatility in our experiment, the variability of results is now significantly larger than in any of the previous instances taken in isolation.

As Figure 13 shows, in this case study replication strategies outperform options-based strategies on average, but with a significant element of risk. The blue dots above the orange line represent replication strategies that performed better than the option-buying strategy, while those below did worse.

Figure 14 Distribution of Outperformance Under Base Assumptions and Relaxed Assumptions (Probability)

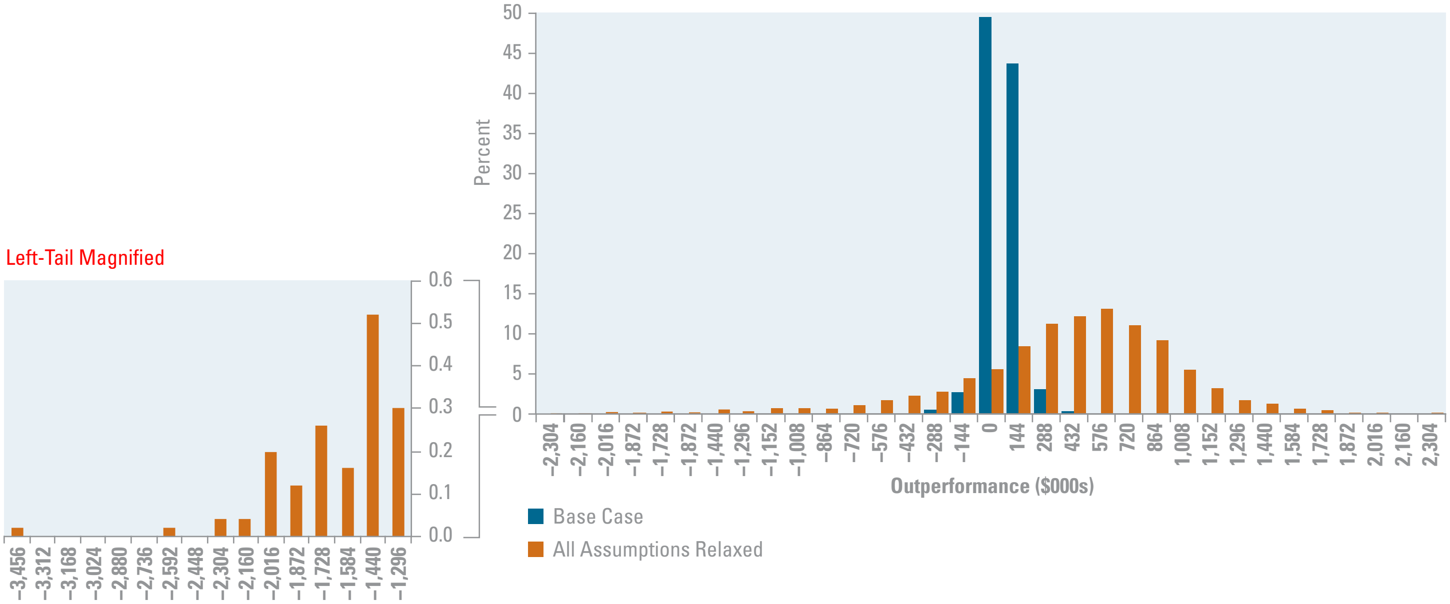

Figure 14 displays the range of out- and underperformance in a Black-Scholes universe (the blue columns), and with all assumptions relaxed (the orange columns), along with their associated probabilities. The small chart displays with more detail the most severe outcomes. In addition to the range of favorable outcomes when all assumptions are relaxed, there is the possibility of significant underperformance — about a 3% chance of underperforming by $1 million or more, and a worst-case underperformance of about $3.5 million.

Also, it's important to recognize that the parameters for this simulation were obtained either from distinct historical periods or from qualitative estimation, and that the techniques used to stress certain Black-Scholes assumptions were highly discretionary. Thus, in practice, the true tail risk can be even larger than the "worst-case" outcomes shown here, particularly in a prolonged recession, severe dislocation or other economic scenario exhibiting great uncertainty or excessive volatility.

Like all great theoretical frameworks, Black-Scholes loses a bit of luster when making the transition from the academic desk to the trading desk, in the sense that the conditions necessary for its application are impossible to achieve in the real world. That being said, it gives a convenient framework within which derivatives markets can operate, and with a little work, an appropriate setting within which to examine tradeoffs associated with both options- and replication-based strategies.

Options-based hedging, in this case study, costs much more on average but exposes the user to far less risk. It's obvious that the overall success of dynamic replication depends on a variety of factors that cannot be known with certainty, all of them outside the Black-Scholes assumptions but within the actual experience of capital markets. It's a decision that requires careful consideration of all factors, and which must be made in concert with the larger portfolio objectives and risk/reward preferences of the investing institution.

▸ Option Pricing via Dynamic Replication

In the example that follows, we posit that a stock now trading at $100 can move to one of two points over the coming year, either upward to $128.40 or downward to $77.88. In addition, we assume that the risk-free rate is 5%, and therefore that $1 invested today in a bank account will be worth $1.05 in a year, irrespective of the stock price. The question now is to determine the fair price for a one-year European put option on this stock with an exercise price of $100. (European put options are exercisable only at expiration — in this case after a year — as opposed to American-style options, exercisable anytime during the option's life.) In exchange for a premium (the cost of the option), our hypothetical put option allows the holder to sell the stock for $100 at maturity even if its market price ends up lower than that. (In practice, it would pay the excess of the exercise price over the stock price.) **Figure 1** diagrams our hypothetical market.

Figure 1

The theory behind dynamic replication is that the payoff of an option is replicable under any market outcome by constantly rebalancing a portfolio of stock and cash so that it always holds Δ shares, where Δ equals the sensitivity of the option's price with respect to the price of the underlying stock. We can find Δ by looking at the ratio of how significantly the value of the option can change vs. how significantly the stock price can change. Assuming the option is issued "at the money" — its exercise price is the stock's current market price — its value can go from zero (if the stock price rises) to $22.12 (if the stock falls). The stock price itself can go from $128.40 to $77.88. Therefore, $\Delta = (0 - 22.12) / (128.40 - 77.88) = -0.4378$. In order to replicate, that's the fraction of a share we'll need to sell short.

To determine how much cash we'll need to invest in the bank account, we need look only at replicating the option's payout in either of the two states. Choosing the "up" state, we know the option payoff is zero. We want the value of our replicating portfolio to equal the payoff of the option; thus, we'd like to solve for *B* such that $0.00 = (-0.4378 x $128.40) + (*B* x 1.05). Solving for *B*, we get $53.54.

It's easy to verify that this replicating portfolio has the same value as the option in the "up" state, since (-0.4378 x $128.40) + ($53.54 x 1.05) = $0.00, and that it has the same value as the option in the "down" state, since (-0.4378 x $77.88) + ($53.54 x 1.05) = $22.12.

Now consider the cost of acquiring this portfolio. We earned $43.78 from the short sale, but invested $53.54 in the bank account. The net cost of establishing this replicating portfolio is therefore $53.54 - $43.78 = $9.76. If we assume that the market admits no arbitrage, then $9.76 must also be the value of the option.

To see why, consider what would happen if the option price were $12. Investors would recognize that they could sell the option for $12, buy the replicating portfolio for $9.76, and earn a risk-free profit of $2.24. After all, since the replicating portfolio produces the same cash flows as the option, the investors would be hedged in all scenarios. They would therefore have a risk-free profit opportunity. We would have an analogous arbitrage opportunity if the option were priced too cheaply (less than $9.76), since investors would sell the replicating portfolio and buy the option, pocketing a risk-free profit. Efficient-market theory dictates that market demand for either trade would force the prices of the option and its replicating portfolio into equilibrium.

Note that nowhere in this exercise did we estimate the probability that the stock price would rise or fall. The only information required of the stock price was how significantly it could rise or fall over the ensuing one-year period. Practitioners call this "volatility." There is a unique option value for a given level of volatility. Likewise, there is only one level of volatility that produces a given option price in concert with the Black-Scholes formula. That level is known as "implied volatility."

To be sure, the previous example is a simplified one. There is only one time horizon and only two possible outcomes, whereas in the real world there would be multiple time horizons and any number of outcomes. However, it is fairly straightforward to generalize this concept by subdividing the one-year interval into smaller time steps. **Figure 2** illustrates the procedure in quarterly steps.

Figure 2

When there is more than one time period, the procedure remains analogous to that of the one-period example. Starting from the option expiration time, we compute the value of the option at maturity simply by using the payoff function max(0,100 – S). After that, we can "roll back" in the tree using at each node the replication methodology described in the one-period example. Reaching the initial node, we find the value of the option. In the case of quarterly time steps, it's $6.92.

Now that we have introduced the concept of replication, let's bridge the gap between replication and the Black-Scholes formula. To better simulate reality, we can continue to shrink the time period between the nodes in our tree, which increases the number of time steps in it. The more we do this, the closer we get to the concept of "continuous trading," in which we perform an infinite number of replication transactions per unit of time. The option value that results from shrinking our time steps to nearly zero (and letting the number of time steps go to infinity) can be interpreted as the Black-Scholes option value.

If the risk-free rate with continuous compounding is r, then for a non-dividend paying stock with initial value S and volatility σ, the Black-Scholes formula gives the value, p, of a European put option with strike price K and time-to-expiration T as

$$p = Ke^{-rT} N(-d_2) - SN(-d_1)$$

$$d_1 = \frac{\ln(S/K) + (r + 0.5\sigma^2)T}{\sigma\sqrt{T}},\ d_2 = d_1 - \sigma\sqrt{T}$$

where $N(\cdot)$ is the cumulative normal distribution function.

To illustrate the connection between continuous rebalancing and the Black-Scholes option value, consider **Figure 3**, which compares the Black-Scholes value of the option to the value obtained through the tree-based replication strategy as the number of time steps increases (and the time between steps decreases). To reiterate, however, both the values from the tree and the value from the Black-Scholes formula are predicated upon an estimate of the volatility – in this case 25%.

Figure 3

Number of Steps in Tree	Time Between Steps	Tree-Based Value	Black Scholes Value
1	1.00000	$9.758	$7.513
2	0.50000	$6.388	$7.513
4	0.25000	$6.920	$7.513
8	0.12500	$7.210	$7.513
16	0.06250	$7.360	$7.513
32	0.03125	$7.436	$7.513
64	0.01563	$7.475	$7.513
128	0.00781	$7.494	$7.513
256	0.00391	$7.504	$7.513
512	0.00195	$7.508	$7.513
1,024	0.00098	$7.511	$7.513
2,048	0.00049	$7.512	$7.513
4,096	0.00024	$7.513	$7.513

Of course, as the main article explains, the equivalence between dynamic-replication strategies and Black-Scholes option prices breaks down in the real world, where the assumptions underlying Black-Scholes tend to fail.

— Nick Mocciolo

55 Farmington Avenue
Hartford, CT 06105
Toll Free: 866.403.4733
www.himco.com