

Forthcoming in the Journal of Derivatives

September 4, 2002

**Assessing the Incremental Value of Option Pricing Theory
Relative to an "Informationally Passive" Benchmark**

by

Stephen Figlewski

Professor of Finance
New York University Stern School of Business
44 West 4th Street, Suite 9-160
New York, NY 10012-1126

212-998-0712

sfiglews@stern.nyu.edu

The author would like to thank Jonathan Goodman, Peter Carr, Sari Carp, Rob Engle and seminar participants at Warwick University, the University of Strathclyde, The Chinese Finance Association, Columbia University, the 12th Annual Derivative Securities Conference, New York University, and Risk 2002 for many valuable comments and suggestions.

Assessing the Incremental Value of Option Pricing Theory Relative to an "Informationally Passive" Benchmark

ABSTRACT

In modern finance, the value of an active investment strategy is measured by comparing its performance against the benchmark of passively holding the market portfolio and the riskless asset. We wish to evaluate the marginal contribution of a theoretical derivatives pricing model in the same way, by comparing its performance against an "informationally passive" alternative model. All rationally priced options must satisfy a number of conditions to rule out profitable static arbitrage. The Black-Scholes model, and others like it, are obtained by assuming an equilibrium in which there are no profitable dynamic arbitrage opportunities either. The passive model we consider incorporates only the fundamental properties of option prices that must hold to avoid static arbitrage, but has no theoretical content beyond that. We review different measures of model performance and apply them to several versions of the Black-Scholes model and our passive model. As with active portfolio management, it turns out to be not that easy for an "active" model to do a lot better than a well designed passive alternative. For example, "classical" Black-Scholes model turns out to be less accurate than the passive benchmark.

At one time, an "active" money manager might have pointed to a record of positive returns on his stock portfolio as evidence that he was doing a good job. But with the advent of portfolio theory and the Capital Asset Pricing Model (CAPM), earning a positive return on average was seen to be an inadequate benchmark for evaluating the manager's performance, because a "passive" strategy of simply buying and holding a market index portfolio, also earns positive returns on average. A manager with special investment skill should at least do better than simply buying the S&P.

At first, professional managers scoffed at the idea that a passive investment strategy could be a viable alternative to active management. But, using this benchmark we have learned that, in fact, most active portfolios do not outperform a passive investment strategy, even though they may make profits in most years and earn a good return on average. There is now broad acceptance that passive investment in an index fund is quite a sound alternative to an active portfolio strategy, and also a recognition among active managers that they need to work hard on such things as holding down costs, given that they will be evaluated relative to a low-cost operational alternative.

Contingent claims valuation, as exemplified by the Black-Scholes (BS) option pricing model, represents another major pillar of modern finance. Option pricing theory has had enormous success both as a theoretical framework and also as a practical investment tool.¹ The BS model now serves as the nearly universal benchmark to which alternative option pricing models are compared.² But, as a benchmark, the BS model is not based on as strong a foundation as the CAPM's passive portfolio strategy. While an investor may easily set up a passive equity portfolio with the same risk exposure (beta) as an actively managed portfolio simply by dividing funds between a market index portfolio and a riskless asset, implementing the dynamic option replication strategy called for by the BS model is much more difficult.

Derivatives pricing models are derived by assuming profitable arbitrage opportunities will be eliminated in equilibrium, which seems to be a very strong principle. It is easy to

¹ We will use the term "Black-Scholes model" broadly in this paper, to include closely related variants for European options, such as the version adjusted for discrete dividends, Merton's continuous-dividend model, the Black model for futures and interest rate options, and the Garman-Kohlhagen model for foreign currency options.

² In an early example, Macbeth and Merville [1979, 1980] considered generalizing the basic BS model to allow the volatility of the underlying stock to vary systematically with its price. The performance of their new Constant Elasticity of Variance model was examined by comparing its ability to match option prices observed in the market, against that of the BS model. Another of many examples was Rubinstein (1985), which examined a number of alternative models and compared them against the standard of the Black-Scholes model. The BS model is taken to be the benchmark, and a new model that does not fit market prices as well as Black-Scholes is normally dropped from further consideration.

prove that option prices must obey a number of well-known constraints, such as put-call parity, that eliminate profitable static arbitrage. If prices violate one of these constraints, a static position can be set up in the present that will lock in an excess return as of option expiration. But to determine a single fair value for a given option, the Black-Scholes model, and others like it, must assume market conditions that rule out profitable dynamic arbitrage opportunities, as well. This requires absence of transactions costs, and other "perfect markets" assumptions. In theory, riskless arbitrage should occur whenever an option's market price deviates from the model value. But in the real world, options arbitrage is inherently risky and costly. Even for a marketmaker, replicating an option's payoff by dynamically rebalancing a hedge portfolio, as dictated by the theory, is not an operational alternative to simply buying the option in the market.³

Derivatives pricing models typically involve highly sophisticated mathematical analysis and very specific assumptions about asset price processes and the market environment. The result may be intellectually satisfying, but rather remote from real world financial markets. How should one judge what the marginal contribution of advanced theorizing is in practical terms?

We propose applying the same approach in evaluating a theoretical option pricing model that we use in judging the performance of an active portfolio manager. We will evaluate the model's marginal contribution by comparing it to a viable alternative model, that is "informationally passive" in that it does not involve the theoretical apparatus that Black-Scholes and other "active" option pricing models require to rule out dynamic arbitrage. We would like to know what might be called the model's "informational alpha," a measure that would represent how much better its performance was relative to this benchmark. How much do we learn from the theory embedded in the BS model, say, beyond what we already know without it?

Exploring that question is the subject of this paper. Our empirical results come from an analysis of S&P 500 index options during a five-year period 1991-1995. This provides a data sample of over 180,000 observations, drawn from one of the most actively traded, and closely analyzed, options markets in the world. We first consider what standard to use in comparing option pricing models. One common approach in model assessment is to look at goodness of fit statistics, like the R^2 in a regression of market prices on model values. We present several models to illustrate how R^2 works as a performance measure for option models. This leads to some useful insights, one of which is that R^2 is not a sufficiently sensitive measure for our purpose. We then consider root mean squared pricing error (RMSE) and show how closely market prices match model values from the classical BS model, and from variations on it, sometimes called "practitioner Black-Scholes," that are widely applied in real-world option trading and market making.

³ Figlewski [1989] and Green and Figlewski [1999] document the risks in option replication in the real world. Leland [1985] shows that rebalancing a hedged option position continuously, as is called for in the theoretical development of the Black-Scholes and similar models, requires trading an infinite number of shares, with infinitely large transactions costs, over the option's lifetime.

We then introduce a new option pricing model that is consistent with the basic properties of rational option prices derived from static arbitrage, but does not require restrictive assumptions about the asset price process or the market environment. This model's ability to match market prices is compared against Black-Scholes in terms of root mean squared error. We examine overall performance of several related variants of each model, as well as performance on different subsets of options separated according to option type (calls versus puts), moneyness, and maturity.

We then extend the comparison to look beyond pricing accuracy. Although theorists focus on a model's ability to compute option fair values, in practice, an option pricing model tends to be used for hedging more than for pricing. It is certainly possible that a model that does not price options very accurately might still be valuable to traders if it performs well in hedge design. We therefore compare the hedging performance of the models in terms of the RMSE of the hedging error in a delta hedge.

Overall, our results indicate that the marginal improvement in option pricing and hedging accuracy that the Black-Scholes model achieves beyond what is available from an informationally passive model is quite limited. We suggest that such a comparison against a passive "null model" is an appropriate test to which any "active" model derived from more extensive theoretical analysis should be subjected. As with active portfolio management, we should judge an active pricing model by its informational alpha. With further investigation, we may well find that for some purposes, a passive model is an adequate representation of option pricing in the market.

II. Data

The data used in the study consist of prices for calls and puts written on the Standard & Poors 500 Stock Index, contemporaneous values for the level of the index, dividend payout on the S&P, and riskless interest rates.

Options

- **Options:** European S&P 500 Index calls and puts.
- **Dates:** January 2, 1991 through Dec. 29, 1995.
- **Option prices:** Midpoint between bid and offer for the last quote of the day. Source: Berkeley Options Data Base.
- **Index level:** S&P index observed simultaneously with option quotes, from Berkeley Options Data Base.
- **Strike prices:** All available strikes.
- **Maturities:** All maturities less than one year.
- **Bad data:** Data points were removed from the sample if
 - Option prices violated a boundary condition (e.g., call price was below intrinsic value),
 - SAS could not compute the implied volatility (typically only very deep in the money very short maturity contracts), or

- there was an obvious error in the prices (e.g., option price greater than the current index value; very few cases).
- **Total observations:** 183,366

Dividends

- A dividend-adjusted index value for each option is constructed by subtracting from the contemporaneously quoted index the present value of dividends actually paid on the S&P 500 index portfolio from the observation date to option expiration.

Interest rate

- 3 month LIBOR, converted to a continuously compounded rate.

It is not obvious what interest rate should be used for option valuation. The model calls for "the" riskless rate. Academic researchers often use rates on US Treasury bills, carefully matching the bill maturity to option expiration. The interest rate for option pricing should reflect the return that could be earned on a very safe alternative investment, but also the cost of funds to options market participants. T-Bill rates tend to be distinctly lower than rates on other money market securities, and they are surely well below the rate at which a trader could borrow to finance a position.⁴ We prefer to use 3 month LIBOR, which is closer to other money market rates. In any case, we do not expect this choice to have an important impact on the results. Under the market conditions of this time period, the mostly short term options we look at are quite insensitive to the interest rate.

Black-Scholes model prices are computed from the standard formula for European options on a stock that pays known discrete dividends over the period to expiration:

$$\begin{aligned} \text{Call} &= S_{\text{adj}} N(d) - PV(X) N(d - \sigma\sqrt{T}) \\ \text{Put} &= S_{\text{adj}} (1 - N(d)) - PV(X) (1 - N(d - \sigma\sqrt{T})) \end{aligned} \tag{1}$$

where: S = index value;
 S_{adj} = S - PV(divs)
 PV(divs) = present value of dividends paid through option expiration;
 X = strike price;

⁴ T-Bills rates are lower than other rates for several reasons. Unlike bank rates, they are tax-exempt at the state and local level. Secondly, there is a special demand for T-Bills to be used as collateral, that may be expected to distort their pricing in the market. Third, they reflect a credit quality above AAA. All of these factors suggest that a trader would be very unlikely to consider holding a hedged option position, with its attendant risks, if its expected return were not appreciably higher than the T-Bill rate. And, of course, no trader can borrow funds at the T-Bill rate.

$$PV(X) = X e^{-rT}$$

r = riskless interest rate;
 T = maturity;
 σ = annual volatility; and
 $d = \frac{\ln(S_{adj}/X) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$

III. Using Regression Analysis to Judge Model Performance

A common way to test a pricing model is simply to run a regression of observed market option prices on the corresponding theoretical values from the model. If the model demonstrates good explanatory power and the fitted regression coefficients have sensible values, the model is judged to be promising. We will begin by presenting results from regression tests on a series of option pricing models.

The regression equation to be fitted is

$$C_{\text{market},j} = a + b C_{\text{model},j} + \epsilon_j \quad (2)$$

where $C_{\text{market},j}$ and $C_{\text{model},j}$ represent the observed market price of option j and the model value, respectively. C may be either a call or a put price. a and b are the regression constant and slope coefficients, respectively, and ϵ_j is the regression residual. If the model gives an unbiased estimate of the market price, the fitted coefficients should have the values $a = 0.0$ and $b = 1.0$. The higher the regression R^2 is, the better the model matches the market.

This very familiar testing strategy entails the assumption that the model's objective is to match market prices. While that sounds like a reasonable goal, especially to a trader, it rules out the possibility that the model could be right and the market could be wrong. That is, if the model gives true values for options, but the market systematically misprices them, equation (2) might not fit very well, even though the model is correct.

Academic financial economists have a great deal of respect for market prices, and have no trouble with the principle that market prices are true option values, given the information that is currently available. Even so, the principle used to derive a valuation equation is not that the model should match market prices, but rather, that the value of the option should equal the cost of replicating its payoff by dynamically trading between the underlying asset and riskless borrowing or lending. Replication cost does not depend on how options are priced in the market, or even whether a market for options exists at all. Thus, the regression in equation (2) is actually a test of the wrong thing; it requires an additional assumption that the market price for every option equals its true expected cost of replication.

When the purpose is to support option trading in actual markets, however, it is more appropriate to treat matching market prices as a necessary property for a model. A model that systematically deviates from market pricing might be a better indicator of true option value than the market but still be of little use to a trader. Traders use observed prices for liquid options to set their bids and offers for less liquid contracts. They then construct hedged positions using those liquid options in order to manage the risk exposures of their positions in the illiquid ones. Since all of these pricing and hedging activities involve options valued at market prices, participants in real world options markets want a model that can match those market prices, almost regardless of whether market prices are "right" or "wrong."

Here are the estimation results for our first model. (t-statistics are given in parentheses.)

Model 1:

$$C_{\text{market}} = 5.637 + 0.955 C_{\text{model 1}} \quad R^2 = 0.9276 \quad (3)$$

(411.3) (1532.7)

This model explains well over 90% of the variance in observed market prices and the t-statistic on the slope coefficient is very large. The coefficient estimates do indicate that the model has some bias, since an unbiased model would have a constant of 0 and a slope coefficient of 1.0, while both of these coefficients are significantly different from those theoretical values. Nevertheless, many researchers might conclude that these results provide strong confirmation of the validity of the model.

But Model 1 is just the option's intrinsic value, $C_{\text{model 1}} = C_{\text{intrinsic}}$, where $C_{\text{intrinsic}}$ is defined by

$$C_{\text{intrinsic}} = \begin{cases} \text{Call: } \text{Max} (0, S_{\text{adj}} - \text{PV}(X)) \\ \text{Put: } \text{Max} (0, \text{PV}(X) - S_{\text{adj}}) \end{cases} \quad (4)$$

Intrinsic value is obviously an important determinant of option value, but it scarcely qualifies as a pricing model. This leads us to several observations.

First, while highly significant coefficients and an R^2 statistic over 0.90 would generally be interpreted as evidence that one has a good model, that conclusion is clearly not appropriate in this case. This shows that one must use caution in interpreting results from the equation (2) regression, if they are presented as showing strong empirical support for a given option pricing model.

Second, intrinsic value alone explains nearly 93 percent of the variance in market option prices. The purpose of a formal option pricing model, therefore, is to explain at most the remaining 7 percent, i.e., the variance of $(C_{\text{market}} - C_{\text{intrinsic}})$. We might even take option

intrinsic value as a very simple kind of passive model, and judge an active model, not in terms of its overall goodness of fit, its R^2 , for example, but in terms of its marginal R^2 , i.e., its improvement over R^2 from using intrinsic value alone.

Third, extending this reasoning, it may make sense to use intrinsic value as a control variate in fitting option model parameters, such as implied volatilities, and in evaluating model accuracy. For example, rather than calculate the implied volatility (IV) from a set of options prices by minimizing

$$\sum_j \left(C_{\text{model},j}(\text{IV}) / C_{\text{market},j} - 1 \right)^2$$

one might minimize

$$\sum_j \left((C_{\text{model},j}(\text{IV}) - C_{\text{intrinsic}}) / (C_{\text{market},j} - C_{\text{intrinsic}}) - 1 \right)^2$$

This will produce the implied volatility that minimizes the model's average discrepancy relative to the market, concentrating only on the portion of the option value that actually depends on volatility.

Let us now look at a valuation model with more economic content. Here are the regression results for Model 2.

Model 2:

$$C_{\text{market}} = 1.786 + 1.012 C_{\text{model 2}} \quad R^2 = 0.9832 \quad (5)$$

(247.0) (3278.8)

This model explains over 98% of the variance of option prices in the market. There is still some bias, since the coefficients still differ significantly from $a = 0.0$ and $b=1.0$, but by most standards, this model would be judged to be highly successful. Even so, Model 2 is considered to be of rather limited value by most academics and almost all traders.

Model 2 is the Black-Scholes model, with the volatility parameter set equal to historical volatility over the previous 250 trading days. The BS equation requires the volatility of the underlying asset from the present through option expiration as an input, but future volatility can not be observed. One way to forecast volatility is simply to assume that future volatility will be the same as realized volatility from a sample of recent price data. There are numerous variations on historical volatility, using different numbers of past observations, calculating a sample mean or constraining it to 0, weighting observations inversely according to their age, etc. Figlewski [1997] and Green and Figlewski [1999]

examine a number of alternative procedures in terms of root mean squared forecast error. Performance varies for the different techniques, but one of the most consistent results is that none of the volatility models seems to provide very accurate forecasts.

The Black-Scholes model with a simple historical estimate as the volatility input is not very useful for trading purposes. The typical situation is that model values for nearly all options will be either below, or above, current market prices. It is not helpful if a trader's model shows that all options are currently underpriced in the market, for instance. A trader wants a model that prices options about the same way the market is pricing them, so that the model can be used to compute bids and offers for off-the-run contracts that will be consistent with the way the market is currently pricing actively traded options.

Academics also distrust Black-Scholes prices computed from historical volatility, because they believe the market makes better volatility forecasts than can be obtained from simple historical data. Among other problems, historical volatility is inherently backward-looking, while traders can be expected to have knowledge about future events that will influence volatility but are not incorporated in a historical estimate. Thus, the volatility input to Model 2 is considered to be inferior to what the market is using.

Model 3 attempts to correct this problem.

Model 3:

$$C_{\text{market}} = 0.057 + 1.008 C_{\text{model 3}} \quad R^2 = 0.9946 \quad (6)$$

(13.3) (5796.5)

Relative to Model 2, Model 3 improves the R^2 by only about 1%, but the fit is extraordinary. There is only about one half of one percent of market variance that the model does not capture.

Unfortunately, in terms of judging model performance, Model 3 is a bit of a cheat. Worse, even though it is a cheat, Model 3 is still felt to be of little use by traders.

Model 3 is the Black-Scholes model using implied volatility. Each day's implied volatility is computed as the value that minimizes squared pricing error, $(C_{\text{market}} - C_{\text{model}}(IV))^2$ across all options with the same maturity. This procedure takes current option prices as an input, then derives the volatility parameter that gives the best overall fit.

The reason Model 3, Black-Scholes with implied volatility, is a cheat is that it is clearly circular reasoning to take current market prices, solve for the specific value of an unknown input parameter that gives the closest fit between model and market, and then judge the model by how close that fit is. To reduce this problem, it is customary for researchers to do "out of sample" tests of their models, by computing implied volatilities

from option prices on day t and then evaluating model performance using those same IV values to price options on day t+1.

Option traders, who are not trying to test models, but to devise consistent price quotes and hedging strategies, are not bothered by lack of independence between market and model prices. They will normally use IVs drawn from contemporaneous market option prices in their models.

Since we are interested in evaluating model performance, not making markets, we rerun the regression for Model 3. Model 4 is Black-Scholes with implied volatility computed from the previous day's market prices.

Model 4:

$$C_{\text{market}} = 0.050 + 1.008 C_{\text{model 4}} \quad R^2 = 0.9944 \quad (7)$$

(11.3) (5678.8)

The fit for Model 4 is almost as good as for the in-sample Model 3. All results we report in the paper from this point on use the same one day out of sample procedure for implied parameters.

We must recognize, however, that a good fit for this model does not allow us to distinguish clearly between the following two hypotheses:

H₀: Model 4 gives good estimates of true option value.

H₁: Model 4 does not give good estimates of true option value, but mispricing in the market is similar from one day to the next. Using yesterday's implied volatility incorporates yesterday's mispricing into today's model values, which produces a good fit to today's market prices.

In other words, Model 4's very close fit to the data does not allow us to conclude that it is a correct model.

Why do traders not find Model 4 useful, even though from the regression results it is hard to believe one could do a lot better? The problem is the volatility "smile" (or, in this market, the "skew"). The original BS derivation treated volatility as a fixed and known parameter for the underlying asset. But empirical observation quickly revealed this to be untrue. The basic model can be easily altered to accommodate volatility that changes over time, so long as it is non-stochastic. With time-varying volatility, the volatility input to the model can differ for different expiration dates, but all options maturing on the same date should be priced with the same volatility. But imposing this constraint in computing IVs from actual S&P option prices leads to systematic mispricing relative to the market as a function of the option's strike price.

Many option markets exhibit a volatility smile, with both in and out of the money options having implied volatility higher than that for the at the money contract. The shape of the smile varies over time, but the basic pattern is a very regular property. S&P 500 index options also exhibited a volatility smile in the first few years they were traded, but after the stock market crash of 1987, the standard pattern in the market changed to a more asymmetric "skew," in which low strike options (out of the money puts and in the money calls) are priced on relatively high volatilities, and IVs fall more or less monotonically as the strike price increases.

When there is a volatility skew, using the same volatility input for all options with the same maturity tends to make low strike options appear consistently too expensive in the market and high strike options look too cheap. Academic researchers take the existence of a smile or skew in the data as evidence that a different pricing model is needed, perhaps one allowing stochastic volatility or discrete jumps in the stock price. Traders tend not to care what the reason for the skew is, but they want a model that fits market option prices for all strikes, so it is standard practice simply to allow different implied volatilities for different options.

Models 5, 6, and 7 reflect different ways of weakening the constraint that IV must be equal for all options with the same maturity (always computing model prices using the previous day's IV estimates, as mentioned above).

Model 5:

$$C_{\text{market}} = 0.054 + 1.008 C_{\text{model 5}} \quad R^2 = 0.9945 \quad (8)$$

(12.3) (5723.7)

Model 6:

$$C_{\text{market}} = -0.017 + 1.002 C_{\text{model 6}} \quad R^2 = 0.9992 \quad (9)$$

(-10.7) (15349.2)

Model 7:

$$C_{\text{market}} = 0.012 + 1.000 C_{\text{model 7}} \quad R^2 = 0.9993 \quad (10)$$

(8.1) (16088.8)

Model 5 allows calls and puts to have different implied volatilities, but they must be equal for all strikes. This still suppresses the skew pattern, but does permit different pricing overall for calls and puts. This small weakening of the constraint makes little difference: the estimation results are virtually identical to those shown for Model 4.

Model 6 allows a smile or skew pattern, with different implied volatilities for different strikes, but requires calls and puts with the same strike to have equal IV. An important reason to expect the latter constraint to hold is that otherwise put-call parity would be violated, leading to the possibility of static arbitrage. The fit of this equation is extraordinary. All but 0.08 percent of the variance of market prices is explained by the

model, the constant is very close to zero and the slope is extremely close to 1.0, even though with such a large sample size, both are still significantly different from the target values.

Finally, Model 7 allows each individual option to have its own implied volatility, producing the best regression fit of all. Practitioners (and academics who really want option model values to match market prices) typically use Model 6 or 7.

The results from this regression analysis lead to two observations.

- The regression test we have been looking at is not really useful for distinguishing a good option model from an unsatisfactory one, because R^2 statistics are all so high.
- Traders force the Black-Scholes equation to match market prices by fitting different implied volatilities for different options on the same underlying. While this allows them to price and hedge illiquid options consistently with current market valuations, the "model" is now inconsistent with the financial theory from which it was derived.

We will discuss these observations in turn.

Our results have shown that even an unsatisfactory model can obtain a high regression R^2 . A more meaningful measure of goodness of fit is the average size of the pricing error produced by the model. The standard measure of this is root mean squared error (RMSE):

$$\text{RMSE} = \sqrt{\sum (C_{\text{market}} - C_{\text{model}})^2 / N} \quad (11)$$

where N is the number of options in the sample.

Exhibit 1 shows the RMSEs and R^2 statistics for the models we have been considering. These results reveal how different these models are in terms of how closely they match market prices, even though they all have extremely high R^2 statistics. Restricting IV to be equal for all options of a given maturity leads to RMSE pricing errors of about \$1.40; allowing a volatility skew with different IVs for different strikes gives RMSE of under \$0.50.

However, since it is not possible for a single underlying asset to have more than one volatility over a specified time period, fitting different IVs for different strike prices means that the valuation equation is in conflict with the underlying theory. Traders tend not to be bothered by theoretical niceties, if they have a tool that works. But what makes the Black-Scholes equation a useful tool for traders, when the way they use it can not be justified by the theory from which the model is created?

We have earlier pointed to two different classes of arbitrage-based relationships for option prices. Those that come from static arbitrage, or more generally from portfolio dominance, are very general and robust, because violation would permit profitable arbitrage on positions that are both (almost) riskless and easily set up in the real world.⁵ Put-call parity is a good example of such a relationship. It must hold for European options, regardless of the behavior of the underlying asset, interest rates, etc.

Theoretical derivatives pricing models, like Black-Scholes, are based on dynamic arbitrage, which is only riskless and costless in theoretical markets, not the real world. While different assumptions about the asset returns process or the market environment lead to different dynamic replication strategies, and therefore different theoretical valuation models, every properly specified pricing model must satisfy the static arbitrage portfolio dominance constraints to be internally consistent. Merton [1973] presents a large number of these option properties and proves that they must hold under general conditions

Thus, a very important feature of the BS equation is that it satisfies all of these constraints. There are quite a few of these relationships. Let us focus on several that are particularly important in this case.

1. European option values go to zero as they get further out of the money.

$$\begin{aligned} \text{Call}(S, X, T) &\rightarrow 0 \quad \text{as } S \rightarrow 0 \\ \text{Put}(S, X, T) &\rightarrow 0 \quad \text{as } S \rightarrow \infty \end{aligned} \tag{12}$$

2. As European options get further in the money, the call (put) value approaches the value of a long (short) forward contract with the same strike price.⁶

Let $F(S_{\text{adj}}, X, T)$ be the forward price to buy, for a price of X on date T , one unit of an underlying asset whose current market price less the value of future dividend payout through expiration is S_{adj} . $PV(X, T)$ is the discounted present value of X dollars to be paid at date T with certainty. Then

$$\begin{aligned} \text{Call}(S, X, T) &\rightarrow F(S_{\text{adj}}, X, T) = S_{\text{adj}} - PV(X, T) \quad \text{as } S \rightarrow \infty \\ \text{Put}(S, X, T) &\rightarrow -F(S_{\text{adj}}, X, T) = PV(X, T) - S_{\text{adj}} \quad \text{as } S \rightarrow 0 \end{aligned} \tag{13}$$

⁵ Portfolio dominance refers to a situation in which one portfolio can never pay off less than a second one, pays off more under some states of the world, and costs less to set up. The first portfolio dominates the second and would be preferred by all traders, so they can not be in equilibrium at current prices.

⁶ The deeper an option goes in the money, the greater is the probability it will be exercised at expiration. In the limit, eventual exercise is certain and the "option" is effectively a forward contract.

3. At the money options have positive time value. We will define "at the money" in this context in terms of the forward, as the value of S such that $S_{adj} = PV(X)$.

$$\begin{aligned} \text{Call}(S_{adj} = PV(X), X, T) &> 0 \\ \text{Put}(S_{adj} = PV(X), X, T) &> 0 \end{aligned} \tag{14}$$

4. Between deep out of the money and deep in the money, the option value is continuous and convex in the asset price and the strike price.⁷

5. Put-call parity holds.

$$\text{Call}(S, X, T) - \text{Put}(S, X, T) = S_{adj} - PV(X) \tag{15}$$

Exhibit 2 plots the familiar call value function to illustrate some of these conditions. Thus, a very important feature of the BS equation is that it satisfies all of these constraints from static arbitrage, that any acceptable option valuation model should obey. Secondly, it has a free parameter, the volatility, that allows the model value to be adjusted to match the market price of an option. This may be enough to make it a valuable tool for a trader, even if the way it is used entails logical inconsistency vis-a-vis the model's underlying theory.

IV. The Informationally Passive Implied G Model

The Black-Scholes equation clearly provides substantial incremental explanatory power relative to the option's intrinsic value, which we looked at above as Model 1, even when historical volatility is used. However the best fit to market prices requires multiple implied volatilities. To judge how much a formal pricing model like Black-Scholes really adds to our understanding of option valuation in the real world, beyond what we already know from the constraints on option prices to eliminate static arbitrage, we would like to evaluate its performance against a viable "passive" option pricing model. This "null model" would serve as a benchmark for evaluating the marginal contribution of a theoretical model in the same way that passive investment in the market portfolio does in assessing the value added by active investment managers and strategies. What we want to know might be called the model's "informational alpha."

⁷ These properties are plausible for options, in general, but in fact, only continuity and convexity with respect to the strike price, given the asset price, is required by portfolio dominance. Continuity and convexity with respect to the stock price require regularity assumptions on the asset price process.

If, in practice, the theoretical underpinnings of the Black-Scholes model are abandoned and the equation is just used as a convenient functional form for extrapolation and interpolation, it is not unreasonable to compare its performance to that of a simple curve fitting procedure that satisfies the portfolio dominance constraints, but has no economic content beyond that.

The following equation has the right shape for a call option's value.⁸

$$y = \sqrt{1 + x^2 / 4} + x / 2 \quad (16)$$

Equation (16) is illustrated in Exhibit 3. Note that as x goes to infinity, y goes to x ; as x goes to minus infinity, y goes to $|x| + x = 0$. At $x = 0$, $y = 1$, so the function is positive "at the money." And it is continuous and convex throughout its range.

To convert this function into an equation for pricing options, we move it to the right location and introduce a parameter G that allows the overall level to be adjusted upward or downward. We can think of G as representing current "general conditions" in the options market, that determine how high option prices are relative to intrinsic values. Instead of implied volatility, this "model" will be based on implied values for G .

Our informationally passive option pricing model is then given by

$$\text{Call}(S, X, T, G, r) = \sqrt{G + \frac{(S_{\text{adj}} - Xe^{-rT})^2}{4}} + \frac{S_{\text{adj}} - Xe^{-rT}}{2} \quad (17)$$

$$\text{Put}(S, X, T, G, r) = \text{Call}(S, X, T, G, r) - (S_{\text{adj}} - Xe^{-rT}) \quad (18)$$

The variables are as defined above. The put value, equation (18), is derived from put-call parity.

Like the BS model, the Implied G model is flexible enough to be fitted with G constrained to be constant for all options with a given maturity, or to differ across subsets of options. Exhibit 4 compares the Black-Scholes variants shown above as Models 4, 5, 6, and 7 against the comparable versions of the passive Implied G model, on the basis of out of sample RMSE, as defined in equation (11).

Exhibit 4 shows that the basic Black-Scholes model with a single implied volatility for all options with the same maturity--the only version consistent with the underlying theory--is

⁸ The origin of this equation is somewhat obscure. Jonathan Goodman at the NYU Courant Institute suggested it to me in a telephone conversation, when I described what I needed the function to do. Peter Carr later pointed out that Duffie (1988) had used a closely related function for a different purpose. Duffie (1988), in turn, credits Stephen Smale for the suggestion.

actually less accurate in terms of out of sample root mean squared pricing error than the passive model. Black-Scholes RMSE is 1.398, versus 1.280 for the Implied G model. Allowing calls and puts to be priced with different values of the implied parameter, but suppressing the smile/skew dependence on option strike price, improves the fit of each model slightly, but the passive model still performs better. Only when the models are relaxed to accommodate the smile does the Black-Scholes equation achieve a lower RMSE than the IG model. For Models 6 and 7, RMSE for the BS equation is below that of the passive model by 2.9 and 3.5 cents (5.4% and 7.0%), respectively.

As we consider these results, it is important to note that the Implied G model has been chosen simply to illustrate the concept of comparing the performance of a pricing model against an informationally passive benchmark. No effort has been made to "tweak" the model in any way to improve its performance. Our object is not to find the best model of this type, only to suggest that model performance should be evaluated against a viable alternative such as this one and to show how such a comparison might be done.

The strong performance of the passive model and the small marginal improvement achieved by the standard Black-Scholes equation, even allowing it to depart from theoretical consistency in order to better match market prices are a surprise, perhaps even a shock. In Exhibit 5, we present RMSE statistics for the different model variants on subsets of the data. Because suppressing the skew but allowing different implied parameters for calls versus puts gave only a minor improvement over the basic single implied parameter formulations, we will drop Model 5 from further consideration.

We report the numbers of observations for each of the three model variants and the RMSE pricing errors for the BS and IG equations in each subsample. As before, all model values are computed using the implied parameter value from the previous day. This accounts for the differing numbers of observations across models. For example, if a given option did not trade on the previous day, it must be eliminated from the sample for Model 7 ("One Implied Value per Option") while it can be included in the Model 4 ("Flat Smile") sample because the common implied parameter value computed from other previous day option prices is available.

The first line in Exhibit 5 duplicates the results shown in Exhibit 4. The next two lines show that both models behave similarly for calls versus puts, with puts being priced a tiny amount more precisely than calls on average. As with the full sample results, the BS model shows worse RMSE than the passive model when a single implied parameter per maturity is allowed, but is somewhat more accurate when it is allowed to incorporate a smile.

The next section of Exhibit 5 examines differences in performance by option maturity. The overall sample contains all options with maturities under a year. Of these, about 57% have maturities of less than three months, for which trading activity is greatest. The table breaks maturities into four categories: under 1 month, 1 to 2 months, 2 to 3 months, and above 3 months. Comparisons between the BS and IG models are similar to what we

have already seen, with the Flat Smile ("classical") version of the BS model doing worse than the passive model, but the less constrained versions doing somewhat better.

Comparing across model versions, we see that the Flat Smile models have substantially larger RMSEs than the less constrained variants, but allowing differences by strike is important for near maturities, while full flexibility--a different implied parameter for each option--does not improve out of sample fit. For example, in the 1-2 month subsample, RMSE was 0.421 for the BS model with a different IV for each strike, but IVs for calls and puts with the same strike required to be the same, while it rose to 0.430 when calls and puts were permitted to differ. This indicates that some of the differences in implied parameters across options do not persist from one day to the next. Fitting a separate implied parameter for each option incorporates that random noise into model values.

Comparing relative model performance across maturities, both looking down each column and across columns, suggests that implied parameter differences across options are due to more than just a smooth smile/skew pattern. With a flat smile, pricing errors for long maturities are about 3 1/2 times larger than for the shortest maturities, while this ratio is more like 1.4 for the least constrained model. Empirically, the smile shape itself is less extreme for longer maturity than for short maturity options, so these results indicate that what is gained by fitting different implied parameters for different options at longer maturities is a greater range of pricing diversity across contracts, not just the ability to match a regular smile or skew.

The last section of Exhibit 5 examines the impact of option moneyness on relative model performance. We look at calls and puts separately, recalling that out of the money calls correspond to low strike prices and out of the money puts to high strike prices. Setting breakpoints appropriately to classify options as out of the money (OTM), at the money (ATM) and in the money (ITM) is more complicated than many researchers recognize. If the underlying asset price is 50 and one is considering one-month options with strikes of 45, 50, and 55, it is easy to split the sample into the three categories. But if one is considering S&P 500 options, when the index is at 500 and there are strikes set 5 points apart, it is inappropriate to classify the 495 strike calls as in the money and the 505 strike calls as out of the money. At normal volatility levels for the index, both of these prices are less than one day's standard deviation away from the current price. All three options are really at the money.

By the same token, a proper definition of moneyness should take option maturity into account. For example, if the underlying asset is at $S_0 = 500$ and its volatility is $\sigma = 15\%$, a range of $(-0.5\sigma$ to $+0.5\sigma)$ corresponds to a range of about $489 < S_T < 511$ for the price at maturity of a one-month option, but approximately $464 < S_T < 539$ at the maturity of a one-year contract. Option moneyness should not be defined simply in terms of the difference between the strike price and the current asset price, without adjustment for the asset's volatility and option maturity.

In Exhibit 5, we have classified options by moneyness as a function of how many standard deviations, in terms of $\sigma\sqrt{T}$, that the strike price is away from the current asset

price, where σ is the implied Black-Scholes volatility for the model variant under consideration and T is the time to maturity for the option. This formulation has the advantage that the probability an option in a given category will end up in the money at expiration is largely independent of volatility or option maturity.

The moneyness categories are as follows:

- Deep OTM -- more than 1.5 standard deviations out of the money;
- OTM -- 1.5 to 0.5 standard deviations out of the money;
- ATM -- 0.5 standard deviations OTM to 0.5 standard deviations ITM;
- ITM -- 0.5 to 1.5 standard deviations in the money;
- Deep ITM -- more than 1.5 standard deviations in the money.

Results for the Flat Smile models show that restricting them to a single implied parameter per expiration has a substantial impact on how well they fit market prices. The worst fit in dollar terms for the BS model is for out of the money options. Deep out of the money contracts also have sizable errors, which are the largest in percentage terms, given the low market prices of these options. The Implied G model, by contrast, fits worst for low strike options (OTM puts and ITM calls), doing much better than Black-Scholes for OTM calls and ITM puts, and much worse for ITM calls and OTM puts.

When multiple implied parameters are fitted, the BS model has lower RMSE than the passive model in nearly every case, with the difference being greatest for the in the money contracts. It is interesting to note that both models do substantially worse for out of the money contracts than in the money contracts.

There are a number of interesting features to be seen in the detail presented in Exhibits 4 and 5, but the general message is that while the performance statistics for the BS model presented in Section 2 appeared very strong, in fact its marginal improvement in matching market prices for options over a well-designed informationally passive model turns out to be modest, at best. Indeed, in order to achieve lower RMSE than the passive model at all, it was necessary to use the Black-Scholes equation in a manner that is inconsistent with the underlying model.

Root mean squared pricing error is an important measure of model performance, but it is clearly not the only thing a user cares about. An obvious alternative would be to measure pricing error in percentage terms rather than in dollars. More relevant, however, is to compare models on the basis of hedging accuracy. It can be argued that in practice, hedging and risk management for options positions is a more important use for a valuation model than pricing is. After all, the market prices for options are treated as fundamental values, to which model prices are adjusted through the use of implied parameters. What traders need the model for is to value options for which there are no good market prices, and especially to understand and manage option risk exposures. We examine the relative performance of the Black-Scholes and the passive models for hedging in the next section.

V. Model Comparisons Based on Hedge Performance

Using a delta neutral hedge to insulate an option position from the effects of stochastic changes in the price of the underlying asset is central to development of modern option pricing theory and its practical application in real world markets. In a perfect markets Black-Scholes world, going long an option and taking an offsetting position in delta shares of the underlying produces a position that is perfectly riskless over the next instant in time. In practice, of course, there are market frictions and positions can not be rebalanced continuously, so a perfect hedge is not feasible. Even so, delta hedging to manage option risk is standard practice, and an important function of an option pricing model is to provide deltas for this purpose. A good hedge is one that minimizes variance, even though it can not be fully eliminated. We now look at performance in hedging options for our Black-Scholes model variants relative to the passive model.

Delta is the partial derivative of the option price with respect to a change in the price of the underlying asset. The exact definition of delta depends on the model and what is assumed about dividends. In our case, the dividend-adjusted stock price is defined as $S_{adj} = S - PV(\text{future dividends})$, so that the dollar amount of dividend payout is assumed not to be affected as the stock price moves. The partial derivative of S_{adj} with respect to S is therefore 1.0. For the Black-Scholes model, this gives

$$\begin{aligned} \text{BS Call delta} &= N[d] \\ \text{BS Put delta} &= N[d] - 1 \end{aligned} \quad (19)$$

where, as above, $N[\cdot]$ is the cumulative normal distribution function and

$$d = \frac{\ln(S_{adj}/X) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$$

Taking partial derivatives with respect to S in equations (17) and (18) gives the deltas for the Implied G model.

$$\text{IG Call delta} = \frac{1}{2} \left(G + \frac{(S_{adj} - Xe^{-rT})^2}{4} \right)^{-\frac{1}{2}} \left(\frac{S_{adj} - Xe^{-rT}}{2} \right) + \frac{1}{2} \quad (20)$$

$$\text{IG Put delta} = \text{IG Call delta} - 1$$

We examined hedge performance by setting up delta neutral hedged positions that were assumed to be either long one call option and short delta "shares" of the S&P 500 index, or long one put and long delta shares. The position was assumed to be held for one day, or until the date of the next available price record in the data set for that option. (Options

were eliminated from the sample if the resulting holding period for the delta hedge would be more than 5 days.) The hedged positions were then unwound at market prices, or if the ending date was option expiration, at the terminal date intrinsic value. Hedging performance is measured by root mean squared deviation from zero change in the position's value. We ignore the one-day expected return on the position, which is very small for a short holding period (e.g., if the riskless rate is 5%, the theoretical expected one-day return on a delta hedge is less than 1.4 basis points). In any case, leaving out the expected value of the hedge return should not affect the comparison between models.

Exhibit 6 presents the comparison between the Black-Scholes and passive models in terms of root mean squared hedging error. Overall, delta hedges based on these models show RMSEs of about half a dollar. Unlike our earlier results, all of the BS models exhibit lower RMSE than their passive counterparts.⁹ Also, despite its distinctly worse performance in pricing, the "classic" Black-Scholes model with a single implied volatility per maturity did just as well as the less constrained variants in hedging. Overall, and for calls and puts separately, the BS models improved on the passive models by around 5% in terms of hedging RMSE.

Breaking the sample down by maturity gave similar results. An interesting feature here is that the performance difference between short and longer maturity contracts was much less in hedging than in pricing.

Turning to comparisons across moneyness, there are both similarities and differences from our previous results. Here again we see little difference in performance for the alternative versions of each model, and the BS models have lower RMSE than the passive models in nearly all cases. However, there are substantial differences in performance between hedges of in the money and out of the money contracts, with the latter showing much higher RMSE. Given that deltas for out of the money options are much closer to zero than for in the money contracts, dollar sizes of hedged positions are smaller for OTM than ITM options, making the differences in dollar RMSEs we see here much larger in percentage terms.¹⁰

The BS model hedges show marginal improvement in performance over the passive model hedge that is small but positive in nearly all cases. Only for deep out of the money puts (and deep in the money calls under one of the model variants) does the passive model beat Black-Scholes in hedging. For in the money contracts the BS performs substantially better than the IG model.

⁹ Notice that taking partial derivatives of (17) and (18) is only one way to obtain an informationally passive model for delta. Better performance in passive hedge design might be obtained from a different formula that was specifically designed to model the behavior of delta.

¹⁰ The difficulty of hedging out of the money options precisely is displayed in detail in Green and Figlewski (1999) for a broad range of markets, maturities, and options.

V. Summary and Conclusion

The CAPM provides a benchmark for evaluating the performance of an active portfolio strategy, by comparing it to a passive portfolio with the same market risk exposure. We have suggested that the same idea of comparing an active strategy against a passive alternative may be usefully applied in evaluating the marginal contribution of a theoretical option pricing model.

We first looked at two common ways of evaluating pricing models. A regression of market option prices on model values provides an R^2 statistic, measuring goodness of fit, and regression coefficients, that indicate model bias. We saw, however, that apparently excellent performance can be achieved by a "model" that is no more than the option's intrinsic value: intrinsic value alone explained more than 92% of the variance in the sample of option market prices. This means that any option model that properly embeds intrinsic value in its valuation should achieve a very high R^2 in the regression, and that the marginal contribution of incorporating more theory is limited to the additional 7+% of variance. The BS model using historical volatility, while largely rejected as a useful valuation tool by both traders and academic researchers, achieved an R^2 of 0.983.

The regression does not discriminate adequately among alternative models whose performance is actually quite different in terms of dollar RMSE. Traders want a pricing model that can match market prices for liquid contracts, while academic researchers want the best estimate of future volatility to put into their pricing equations. The solution in both cases is to use implied volatility. However, stock index option prices in the market do not exhibit the same implied volatility for all options with the same expiration date, contrary to what is required by the underlying theory. The standard practice among traders is simply to fit different IVs for options with different strike prices. But this turns the BS equation from a pricing model that has been rigorously derived from theoretical principles into just a functional form that is useful for interpolation and extrapolation in a trading context.

Rational option prices must exclude profitable static arbitrage. Theoretical values from the Black-Scholes model and all similar option pricing models do so, but to arrive at a unique option fair price, they assume all dynamic arbitrage opportunities are also eliminated. Static arbitrage is easy to implement in the real world, but dynamic replication of option payoffs is not. Even so, the BS model is widely used by traders. But the way it is implemented, as "practitioner Black-Scholes," uses implied volatility extracted from market option prices, with different IVs for different options on the same underlying. In other words, practitioners use the Black-Scholes equation, but not the Black-Scholes model.

This raises the question of how much of the BS model's practical success is simply due to the fact that BS values satisfy the static arbitrage constraints, and how much comes from the real theoretical content of the model, which is essentially the expression for the cost of replicating the option's payoff under the model's assumptions (or equivalently, the cost

of hedging the option dynamically through expiration). We addressed this question by comparing the BS model against an "informationally passive" model, that was just a simple mathematical function which also obeys the static arbitrage constraints. The Implied G model is one such function. We found that the regression and RMSE results looked extraordinarily good for the BS model, but the marginal improvement over the passive model was rather limited. In fact, the classic BS model with a single implied volatility for all options of a given maturity was less accurate than the passive model. When the constraint on IVs was relaxed to accommodate the volatility skew, the BS equation performed better than the passive model in terms of root mean squared pricing error, but only by a few cents.

The relative performance of the BS models, including the classic model, was better when hedging accuracy was the criterion. Interestingly, fitting multiple implied volatilities to allow the skew did not seem to make much difference to model performance in delta hedging.

As researchers, we tend to think that our intellectual creations are of great value, and we are perhaps insufficiently skeptical when empirical tests seem to confirm that belief. What we have seen in this paper is that some kinds of tests are more stringent than others (e.g., it was easier to achieve a high R^2 than a low RMSE), and that if we want a measure of a model's true contribution, we should focus on incremental explanatory power relative to a benchmark that properly incorporates all of the properties of the data that must hold independent of our model. The passive Implied G model is an example of such a benchmark for option valuation.

I believe that subjecting models to this kind of test will help us to evaluate them more honestly and to focus more clearly on how much improvement can be obtained, at what cost, by more ambitious theorizing.

One also may consider whether for some purposes, use of a robust passive model may be preferred to a more ambitious option model that can be harder to work with and require more detailed input information. The IG model also offers flexibility as a pricing function, in that it can be readily extended to incorporate constraints on G across options or additional factors.¹¹ For example, one might want to parameterize the skew pattern by constraining G to be a specific function of moneyness, such as a quadratic, or to add structure across option maturities to capture an empirical "term structure of Gs."

As always, further research in this area would be worthwhile.

¹¹ In giving presentations of this paper, I have occasionally succumbed to the somewhat whimsical and self-referential urge to refer to this as the "Flexible Implied G" model, or the FIG model, for short.

References

Black, Fischer and Myron Scholes. "The Pricing of Options and Corporate Liabilities." Journal of Political Economy, v81, May-June 1973. 1973.

Duffie, Darrell. "An Extension of the Black-Scholes Model of Security Valuation." Journal of Economic Theory, vol. 46 (1988), 194-204.

Figlewski, Stephen. "Options Arbitrage in Imperfect Markets." Journal of Finance, Vol 44, No. 5, Dec. 1989.

Figlewski, Stephen. "Forecasting Volatility." Financial Markets, Institutions, and Instruments Vol.6, No.1. Boston MA: Blackwell Publishers, 1997.

Green, T Clifton and Stephen Figlewski. "Market risk and model risk for a financial institution writing options." Journal of Finance, 1999, v54(4), 1465-1499

Leland, Hayne. "Option Pricing and Replication with Transactions Costs." Journal of Finance, 1985, v40(5), 1283-1301.

MacBeth, James D. and Larry J. Merville. " An Empirical Examination Of The Black-Scholes Call Option Pricing Model." Journal of Finance, 1979, v34(5) , 1173-1186

MacBeth, James C. and Larry J. Merville. " Tests Of The Black-Scholes And Cox Call Option Valuation Models." Journal of Finance, 1980, v35(2) , 285-303.

Merton, Robert C. " Theory of Rational Option Pricing." Bell Journal of Economics and Management Science, 1973, 4 (Spring 1973), 141-183.

Rubinstein, Mark. " Nonparametric Tests Of Alternative Option Pricing Models Using All Reported Trades And Quotes On The 30 Most Active CBOE Option Classes From August 23, 1976 Through August 31, 1978." Journal of Finance, 1985, v40(2) , 455-480.

Exhibit 1: Regression R² and Root Mean Squared Pricing Error for Models 1 through 7

R² statistics come from estimates of regression equation (2), as reported in the text. Root mean squared error is defined as

$$\text{RMSE} = \sqrt{\sum (C_{\text{market}} - C_{\text{model}})^2 / N} \quad (11)$$

Model	R ²	RMSE
1. Intrinsic value	0.9276	7.151
2. Black-Scholes, historical volatility	0.9832	3.097
3. Black-Scholes, implied volatility, same day	0.9946	1.375
4. Black-Scholes, previous day, single IV per maturity	0.9944	1.398
5. Black-Scholes, previous day, put and call IVs differ	0.9945	1.386
6. Black-Scholes, previous day, IVs differ by strike	0.9992	0.510
7. Black-Scholes, previous day, IVs differ for each option	0.9993	0.486

Exhibit 4: Mean Squared Pricing Error for Black-Scholes versus the Implied G Model

The Table presents Root Mean Squared Error in dollars, as defined in equation (11), for the Black-Scholes equation and the Implied G Model, defined in equations (17) and (18). Implied volatility and implied G are computed from previous day option prices.

Model Variant	Black-Scholes	Implied G Model
4. Single implied parameter per maturity for all options	1.398	1.280
5. One implied parameter for puts, one for calls	1.386	1.271
6. Implied parameters differ by strike; calls = puts	0.510	0.539
7. One implied parameter for each option	0.466	0.501

Exhibit 5: Root Mean Squared Pricing Error for Black-Scholes vs. Implied G Model

The Table presents Root Mean Squared Error in dollars, as defined in equation (11), for the Black-Scholes equation (BS) and the Implied G Model (IG) on various subsets of the data sample. Implied volatility and implied G are computed from previous day option prices. Option moneyness is defined in terms of the standard deviation of the log return for the maturity of the option $v = \sigma\sqrt{T}$, as described in the text. Breakpoints for calls are as follows. Deep OTM: $\ln(S_{adj}/PV(X)) < -1.5 v$; OTM: $-1.5 v \leq \ln(S_{adj}/PV(X)) < -0.5 v$; ATM: $-0.5 v \leq \ln(S_{adj}/PV(X)) < 0.5 v$; ITM: $0.5 v \leq \ln(S_{adj}/PV(X)) < 1.5 v$; Deep ITM: $1.5 v \leq \ln(S_{adj}/PV(X))$. Breakpoints for puts are the same, with signs reversed.

Subsample	Flat Smile			One Implied Value per Strike			One Implied Value per Option		
	# obs	BS	IG	# obs	BS	IG	# obs	BS	IG
Full Sample	182,491	1.398	1.280	180,713	0.510	0.539	176,703	0.466	0.501
Calls	92,530	1.390	1.323	91,602	0.515	0.542	89,813	0.492	0.524
Puts	89,961	1.407	1.234	89,111	0.504	0.536	86,890	0.438	0.476
Maturity									
0-1 months	36,258	0.522	0.492	36,110	0.374	0.400	33,742	0.373	0.407
1-2 months	39,021	0.858	0.812	38,677	0.421	0.456	37,776	0.430	0.467
2-3 months	29,543	1.127	1.054	29,032	0.451	0.480	28,714	0.451	0.482
3-12 months	77,669	1.901	1.727	76,894	0.616	0.645	76,741	0.523	0.559

Exhibit 5 continued: Root Mean Squared Pricing Error for Black-Scholes vs. Implied G Model

Subsample	Flat Smile			One Implied Value per Strike			One Implied Value per Option		
	# obs	BS	IG	# obs	BS	IG	# obs	BS	IG
Moneyness - Calls									
deep OTM	13,196	1.105	0.696	6,098	0.442	0.440	4,663	0.531	0.590
OTM	22,293	1.852	0.936	27,738	0.578	0.601	27,390	0.633	0.669
ATM	29,734	1.242	1.195	29,348	0.576	0.596	29,364	0.528	0.551
ITM	17,862	1.494	2.132	16,824	0.453	0.507	15,868	0.274	0.339
deep ITM	9,534	0.250	1.159	11,594	0.243	0.284	12,528	0.071	0.072
Moneyness - Puts									
deep OTM	5,222	0.766	0.865	6,149	0.548	0.529	3,855	0.602	0.649
OTM	15,997	1.818	2.267	16,287	0.670	0.688	16,352	0.664	0.693
ATM	29,769	1.240	1.089	29,338	0.554	0.580	29,331	0.471	0.501
ITM	22,398	1.693	0.712	28,281	0.384	0.453	27,948	0.234	0.314
deep ITM	16,575	0.857	0.505	9,056	0.183	0.231	9,404	0.079	0.099

Exhibit 6: Root Mean Squared Hedging Error for Black-Scholes vs. Implied G Model

The Table presents Root Mean Squared Error in dollars as the change in the value of a delta hedged position that is long one option and short delta "shares" of the index, for the Black-Scholes equation (BS) and the Implied G Model (IG) on subsets of the data sample. Implied volatility and implied G are computed from previous day option prices. Option moneyness is defined in terms of the standard deviation of the log return for the maturity of the option $v = \sigma\sqrt{T}$, as described in the text. Breakpoints for calls are as follows. Deep OTM: $\ln(S_{adj}/PV(X)) < -1.5 v$; OTM: $-1.5 v \leq \ln(S_{adj}/PV(X)) < -0.5 v$; ATM: $-0.5 v \leq \ln(S_{adj}/PV(X)) < 0.5 v$; ITM: $0.5 v \leq \ln(S_{adj}/PV(X)) < 1.5 v$; Deep ITM: $1.5 v \leq \ln(S_{adj}/PV(X))$. Breakpoints for puts are the same, with signs reversed.

Subsample	# obs	Flat Smile		One Value per Strike		One Value per Option	
		IV	IG	IV	IG	IV	IG
Full Sample	181,718	0.495	0.512	0.486	0.514	0.487	0.514
Calls	92,305	0.519	0.536	0.510	0.537	0.508	0.537
Puts	89,413	0.469	0.486	0.461	0.488	0.465	0.489
Maturity							
0-1 months	35,541	0.440	0.450	0.440	0.450	0.443	0.451
1-2 months	38,783	0.468	0.484	0.456	0.482	0.457	0.483
2-3 months	29,834	0.478	0.491	0.465	0.490	0.466	0.491
3-12 months	77,560	0.537	0.559	0.528	0.562	0.528	0.563

Exhibit 6 continued: Root Mean Squared Hedging Error for Black-Scholes vs. Implied G Model

Subsample	# obs	Flat Smile		One Value per Strike		One Value per Option	
		IV	IG	IV	IG	IV	IG
Moneyness - Calls							
deep OTM	14,913	0.637	0.637	0.626	0.638	0.625	0.637
OTM	21,814	0.637	0.682	0.627	0.671	0.628	0.672
ATM	26,781	0.540	0.559	0.538	0.556	0.538	0.556
ITM	18,112	0.332	0.323	0.299	0.361	0.295	0.361
deep ITM	10,685	0.133	0.103	0.124	0.102	0.090	0.100
Moneyness - Puts							
deep OTM	5,626	0.687	0.672	0.690	0.666	0.712	0.670
OTM	16,453	0.663	0.663	0.651	0.687	0.660	0.691
ATM	26,817	0.500	0.517	0.498	0.514	0.499	0.515
ITM	21,906	0.328	0.398	0.313	0.380	0.311	0.378
deep ITM	18,611	0.195	0.197	0.173	0.202	0.167	0.201

EXHIBIT 2: VALUE OF A CALL OPTION

Initial asset = 100, Strike = 100, Maturity = 1 Month, Interest = 8.00%, Volatility = 0.25

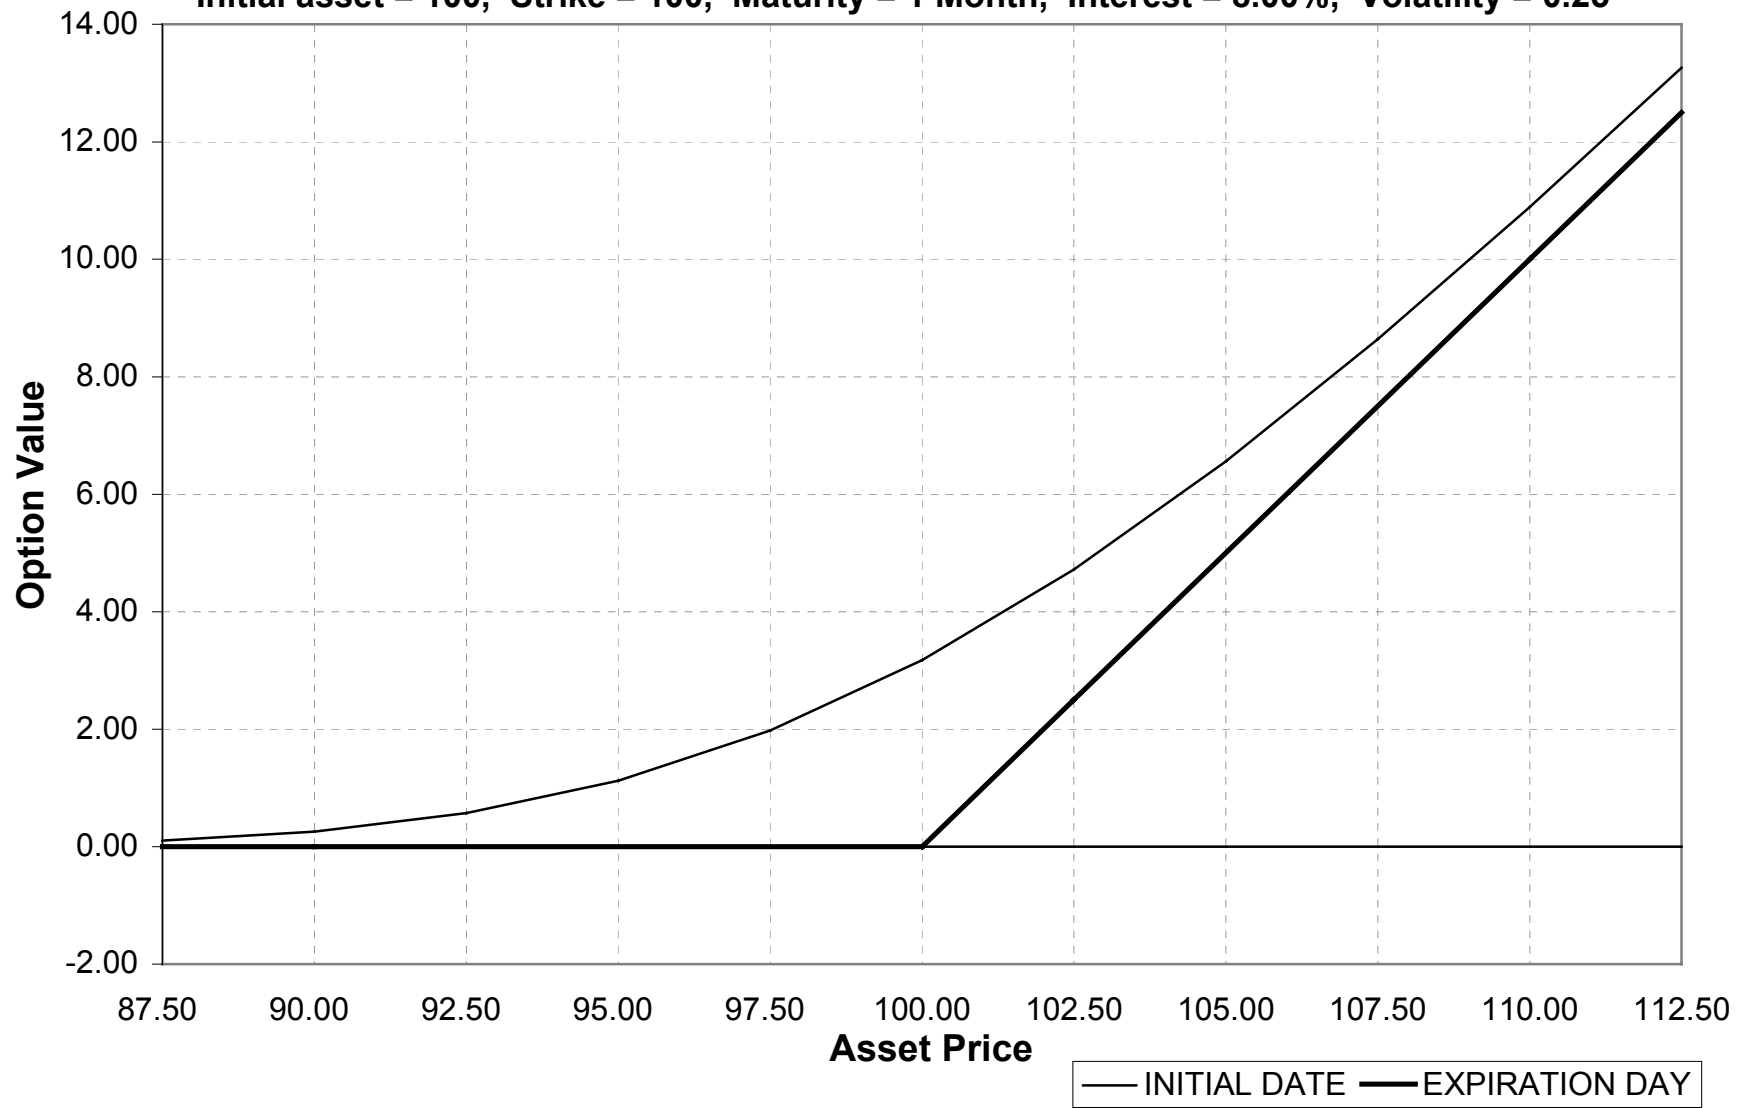


Exhibit 3: Graph of Equation (16)

