

Analyzing Investment Data Using
Conditional Probabilities: The
Implications for Investment Forecasts,
Stock Option Pricing, Risk Premia,
and CAPM Beta Calculations

By

Richard R. Joss, Ph.D.

Resource Actuary, Retired
Watson Wyatt Worldwide

Reprinted from:

Journal of Financial and Economic Practice

Vol. 6 No. 2 (Spring 2006)

**Analyzing Investment Data Using
Conditional Probabilities: The
Implications for Investment Forecasts,
Stock Option Pricing, Risk Premia,
and CAPM Beta Calculations**

ABSTRACT

This paper demonstrates that using a conditional probability approach to analyzing investment data yields investment forecast and stock option pricing models which more closely match actual market data, and it demonstrates the impact that conditional probabilities could have on CAPM beta and risk premia calculations.

It has been a relatively standard practice in financial analysis to assume that observed market returns may be considered to be an independently determined data set for the purpose of statistical analysis. Even if this assumption is not stated explicitly, it may be implied by the mathematical basis for the statistical analysis tool being used. This paper questions this independently determined assumption for historical returns, and offers a new assumption based on the mathematics of conditional probabilities.

Using the new approach, the paper develops an improved investment forecast model and an improved stock option pricing model. Both of the new models are shown to be much more accurate than previous models by comparing the model output with empirical data. The issue is important because of the increased use of both investment forecasts and stock option pricing models.

Investment forecast modeling is increasing in importance as more and more workers rely on defined contribution accounts as the primary source of retirement income. Forecast scenarios help workers decide on the investment strategy and the savings level necessary to meet retirement income objectives. Investment forecasts of expected account growth are also likely to play a significant role in the public policy debate about allocating a portion of Social Security contributions to privately-managed accounts.

There is increased interest in stock option pricing models, and the Black-Scholes stock option pricing model in particular, since the Financial Accounting Standards Board (FASB) issued FAS 123(R) Share-Based Payment in December, 2004. This revised accounting standard specifically mentions the Black-Scholes model as a tool for valuing stock option awards for disclosure purposes in corporate financial statements.¹

Common investment forecasting models and the Black-Scholes stock option pricing model are linked by the fact that they both rely on the assumption that investment returns are lognormally distributed, and that observations of historical data constitute an independently determined set of data. It is the goal of this paper to replace this second assumption with a more accurate conditional probability approach.

This paper is presented in 7 sections, each of which is briefly summarized as follows:

- | | |
|--------------|--|
| Section I: | Lognormal Distributions is used to provide background on this interesting family of probability density functions. |
| Section II: | Parameter Determination – Independent Historical Return Method presents a lognormal parameter selection process that is commonly used at the current time. |
| Section III: | Parameter Determination – Conditional Probability Method presents an alternative parameter determination process. |

- Section IV: A Second Example expands on the illustrations provided in Sections II and III which are based on large company stock returns. This section provides illustrations based on small company stock returns.
- Section V: Stock Option Pricing presents the details of a new stock option pricing model based on the conditional probability method of parameter determination.
- Section VI: Investor Impact presents a brief discussion of the impact that ideas presented in this paper could have on investor behavior, with emphasis on risk premia and CAPM beta calculations.
- Section VII: Summary presents a brief summary.

I. LOGNORMAL DISTRIBUTIONS

Investment forecasts provide investors with information about potential risks and rewards of a particular investment opportunity. In making a forecast, it is common to use historical data as a guide in selecting key forecasting variables. While past performance is no guarantee of what the future holds, with appropriate adjustment for current market conditions this historical guide does provide the user of the forecast with reasonable information which may be helpful in making investment decisions.

Investment return forecasting models found in the literature² are designed to provide not just a single point estimate of future return, but a complete distribution of possible values. If it were assumed that past returns were distributed normally (fit the standard bell-shaped curve), one could use historical data to estimate the mean and standard deviation of a prospective distribution. Having this distribution would then provide answers for questions not only about the expected mean return, but about the probability that future results may exceed, or fall below, any given threshold.

It is a fact, however, that historical results have not tended to fit the traditional normal curve pattern, and a slightly different model is commonly used, the lognormal distribution model. This model is based on the fact that history has shown return relatives on stocks and other economic measures to be skewed toward the larger values. This is expected, since by the very definition of return relatives they can become quite large, but never fall below the value of zero. To compensate for this skewing, the model assumes that investment returns are lognormally distributed. This means that if r is a random variable representing rates of return measured over some fixed time frame, $1 + r$ is its return relative, and the random variable $R = \ln(1 + r)$ is assumed to be normally distributed.

This normal distribution of logarithms is completely determined by its mean (μ) and its standard deviation (σ). Once these parameters are determined, the lognormal forecasting process may be used to provide complete distributions for all investment horizons. To show the development of these forecasts, one needs only to analyze a general lognormal density function, which is defined as follows:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \cdot \frac{1}{x} \cdot e^{-\frac{(\ln x - \mu)^2}{2\sigma^2}} \quad (1)$$

The mathematical expectation, variance, and standard deviation of the random variable x are calculated to be as follows:

$$\text{Expected Value: } E(x) = e^{\mu + \frac{\sigma^2}{2}} \quad (2)$$

$$\text{Variance: } \sigma^2(x) = [E(x)]^2 \cdot (e^{\sigma^2} - 1) \quad (3)$$

$$\text{Standard Deviation: } \sigma(x) = [E(x)] \cdot \sqrt{e^{\sigma^2} - 1} \quad (4)$$

The distribution of wealth after n time periods is given by the equation:

$$W(n)_z = e^{nu + z\sigma\sqrt{n}} \quad (5)$$

where z is the z-score of the percentile in question.

For any given time horizon, n, the wealth distribution described by equation (5) is also lognormal, and has the same density function as shown in equation (1), except that the parameter μ is replaced with $n\mu$, and the parameter σ^2 is replaced with $n\sigma^2$. The expected wealth after n time periods is:

$$EW(n) = e^{nu + \frac{n\sigma^2}{2}} \quad (6)$$

Once the parameters μ and σ^2 are determined, the investor can use the lognormal distribution tool to answer questions such as: What is the likelihood that a \$1,000 investment will grow to exceed \$10,000 after a 15 year investment period?

The only remaining issue is how to use the historical data to determine the parameters μ and σ^2 . While addressing that issue it is helpful to keep in mind the outcomes if the parameters, and the resulting lognormal model, fit ideal circumstances:

- 1) Different approaches to using the same historical data should produce substantially equivalent results. Thus, the forecast should depend on the data itself, and not on how a given forecaster chooses to use the data.
- 2) Long-term wealth projected by the model should be reasonably close to the long-term wealth that serves as the data source. If the long-term expected wealth from the model significantly exceeds or falls below the actual wealth, the investor would be receiving information on potential wealth distributions under the potentially misleading assumption that the future expected returns are either significantly better or significantly worse than what has actually been observed.
- 3) The expected distribution of wealth at shorter-term time horizons should bear a reasonable relationship to the actual distribution of historical results. After all, long-term wealth can only be achieved by the compounding of shorter-term values. Hence, the accuracy of the short-term distribution is essential for accurate projections of long-term wealth.

II. PARAMETER DETERMINATION – INDEPENDENT HISTORICAL RETURN METHOD

The most common method of parameter determination is the one described in Ibbotson (2002). This method uses the arithmetic mean of a historical sampling of data, and the sample standard deviation from the same data. Letting these values be $E_A(r)$ and S_A respectively, the equations for μ and σ are defined as follows:

$$\mu = \ln(1 + E_A(r)) - \frac{\sigma^2}{2} \quad (7)$$

$$\sigma = \sqrt{\ln\left(1 + \left(\frac{S_A}{1 + E_A(r)}\right)^2\right)} \quad (8)$$

These results were derived by simply setting the sample mean equal to the calculated mean:

$$E_A(1+r) = 1 + E_A(r) = e^{\mu + \frac{\sigma^2}{2}} \quad (9)$$

And by setting the sample variance equal to the calculated variance:

$$S_A^2 = (1 + E_A(r))^2 \cdot (e^{\sigma^2} - 1) \quad (10)$$

Equation (10) can be solved for σ , and equation (9) can be solved for μ . The solutions produce the exact values for μ and σ that were presented in equations (7) and (8) above.

To illustrate how these concepts are applied in practice, historical returns from the Ibbotson Associates Stocks, Bonds, Bills and Inflation 2002 Yearbook will be used. As noted above, it is common to adjust probabilistic forecasts for current market conditions. This adjustment might be based, for example, on the difference between the current yield on low-risk bonds, and the corresponding historical yields. Since this paper is concerned with methodology, rather than any specific result, for purposes of this paper the adjustment is assumed to be zero.

Using annual calendar year returns for large company stock, and the formulas (7) and (8) above, the annual parameter μ is determined to be .1033 and the annual parameter σ^2 is calculated to be .03136. However, the annual parameters could also be determined by using monthly large company stock returns and the formulas (7) and (8) above to first determine monthly parameters, and then multiply these parameters by 12 to determine annual parameters. Using this approach, the annual parameter μ becomes .1018 and annual parameter σ^2 is calculated to be .03731. This difference is, of course, just a natural consequence of using actual data. Thus, even though there is a formula provided to determine the parameters, some human judgment is still involved in making an actual forecast; i.e. whether to use monthly data results multiplied by 12, or to use the directly determined annual results.

Unfortunately, these different approaches yield significantly different results. For example the expected annual return is 12.65% when the annual calendar year data is used and 12.80% when the monthly data is used; and it is possible to produce expected returns ranging from 10.8% to 13.8% just depending upon how this single 76-year large company stock return history, and formulas (7) and (8), are used to determine the parameters μ and σ^2 .

Thus, this parameter determination process fails the first of the ideal criteria listed in the previous section. The results from any given forecast depend heavily upon how the forecaster chooses to use the data.³

Also, unfortunately, for the most common ways in which the parameters are actually determined, the model fails the other two ideal criteria as well. The 76-year expected wealth is \$8,584 for the parameters determined from annual calendar year data, and \$9,458 for the parameters determined from monthly data. Both of these results are significantly larger than the actual 76-year wealth for large company stock of \$2,279.

Since the long-term results are so far off, it is not surprising that the short-term expected distribution of results is also significantly different from the actual results. Table I compares the distribution of expected monthly returns using the above parameter determination process with the actual distribution of returns. The expected distribution clearly has a much larger standard deviation than is present in the actual results.

The independent historical return method of parameter determination is based on a philosophy that each calculated historical return, whether it be a monthly, quarterly, or annual return, constitutes an independently determined result. Under this philosophy, the arithmetic mean of these historical results is, in fact, the best estimate for a future return. When these shorter period returns are combined to produce longer period returns through the creation of a lognormal model it is possible to develop a distribution of long-term returns. In this case the actual long-term history which serves as the source for the data winds up near the median of the expected long-term distribution and this median value is less than the mean of the distribution, sometimes by significant amounts. In short, this model fails all three of the ideal criteria, and calls into question the basic philosophy which drives the model.

The next section of this paper introduces a parameter determination process based on a different philosophy. This new process generates a model where the outcome is much more in line with the ideal criteria presented in the prior section.

III. PARAMETER DETERMINATION – CONDITIONAL PROBABILITY METHOD

The parameter determination process in the previous section was driven by the assumption that each historical result constitutes an independently determined event. In this section, a different parameter selection method is developed, one which is based on the assumption that the single observed long-term result should lie at the mean, not the median, of the expected long-term lognormal distribution. The monthly, quarterly or annual returns become just periodic observations of this single long-term result.

Under this method of parameter determination, periodic historical data is used only for the purpose of determining the relationship between the parameters μ and σ^2 using equations (7) and (8) above. Once the relationship is determined, the entire long-term history is used for the purpose of calculating the exact parameter values. While this process still depends on the judgment of the forecaster for determining the relationship between parameters, the use of the entire long-term history for calculating final parameter values significantly reduces the differences between alternative forecasts derived from the same data.

For example, using the data above for the 76-year large company stock returns, μ was calculated above to be .1033 and σ^2 was calculated above to be .03176; thus $\sigma^2 = (.30745) \mu$. A forecaster could decide that this relationship is appropriate, and once this relationship is established, the specific parameters are determined using the following equations; based on the fact that actual 76-year wealth was \$2,279.

$$2,279 = e^{76\mu + 76\frac{\sigma^2}{2}} \quad (11)$$

$$\sigma^2 = (.30745) \mu \quad (12)$$

Solving these two equations yields $\mu = .0882$ and $\sigma^2 = .02712$. Under the conditional probability method of parameter determination, future large company stock returns are assumed to be lognormally distributed, with parameters determined by the above process. Obviously, different data sets will produce different forecasts, but as noted above, the level of variability between different forecasts is significantly reduced.

Also, the parameter determination process presented in this section meets the second ideal criteria mentioned in Section I by its very design. Using this process, the expected long-term wealth for the model will always be the long-term wealth actually observed.

Finally, the distribution of actual historical monthly results needs to be compared with the expected distribution, as determined by the lognormal density function with parameters based on the above determination process. However, the philosophy for the revised method is that the actual long-term historical result is just a single data element from the complete distribution of 76-year returns, and this single data element lies near the mean of expected 76-year results. Hence, to develop an expected distribution of observed periodic results, one must first determine the

conditional probability density function for the distribution of periodic results described by equation (1) given that ending wealth is at its expected value. This density function is given by the following equation:

$$c(x) = \frac{1}{x} \cdot \frac{\sqrt{n}}{\sqrt{n-1}} \cdot \frac{1}{\sigma\sqrt{2\pi}} e^{\left[\frac{-n(\ln x - (\mu + \frac{\sigma^2}{2}))^2}{2(n-1)\sigma^2} \right]} \quad (13)$$

Table II compares the distribution of the monthly results using this density function with the actual 912 calendar month returns for large company stocks, using $n = 912$, $\mu = .00735$, and $\sigma^2 = .00226$ in equation (13). These values of μ and σ^2 are the monthly equivalent values of the annual parameters determined above.

The difference between Tables I and II is quite dramatic. The total difference between actual and expected returns is 256 for the independent historical return method; whereas the total difference drops to 127 for the conditional probability method and this difference could be reduced further if the relationship between μ and σ^2 were reduced to perhaps $\sigma^2 = .26\mu$ or $\sigma^2 = .27\mu$. Furthermore, as noted above, a quick visual scan of Table I shows the expected standard deviation is clearly larger than what was observed in the actual history of results, whereas Table II results are much more consistent with the actual historical data.

This observation is not at all surprising. The independent historical return method yields expected future returns at the arithmetic mean of actually observed historical data, whereas the conditional probability method yields expected future returns equal to the geometric mean of the actually observed data. This distinction is critical since the difference between the geometric mean and arithmetic mean is a measure of the variability of the data. If the data is uniform, the geometric mean and arithmetic mean produce the same result. If the data is not uniform, the arithmetic mean will exceed the geometric mean; and the greater the variability, the greater the difference between the means. If a forecast is based on the arithmetic mean, it counts this variability twice. The variability is fully reflected one time by the geometric mean, and then reflected a second time by the difference between the geometric and arithmetic means.

While this statement seems intuitively obvious from general reasoning, it can also be demonstrated mathematically. The lognormal model yields expected wealth after n time periods as described by equation (6), which in turn yields a

constant expected geometric mean return of $e^{\mu + \frac{\sigma^2}{2}}$ for all time periods. However, the expected value of the conditional probability density function in equation (13) is:

$$E_c(x) = e^{\mu + \frac{\sigma^2}{2} + \frac{(n-1)\sigma^2}{2n}} \quad (14)$$

Since the underlying returns vary from year to year, the above equation (14) shows that historically observed arithmetic means contain an extra variability component if the expected wealth described by the lognormal process is actually achieved. These larger arithmetic means do not affect the expected value of wealth which is always given by equation (6).

In summary, annual parameters $\mu = .0882$ and $\sigma^2 = .02712$ used as part of a lognormal forecast process for large company stock returns:

1. Produce a distribution of results which matches historical data reasonably well. (See Table II.)
2. Yield an expected 76-year wealth of:

$$e^{\left[76(.0882)+76\left(\frac{.02712}{2}\right)\right]} = \$2,279 \quad (15)$$

3. Yield an expected 76-year arithmetic mean of:

$$e^{.0882 + \frac{.02712}{2} + \frac{(75)(.02712)}{(76)(2)}} = 12.2\% \quad (16)$$

This last result compares favorably with an actually observed arithmetic mean rate of return which varies from 10.8% to 13.8% depending upon how the data is used to determine the arithmetic mean.

IV. A SECOND EXAMPLE

Based on the equations above, one would expect that the more variable the underlying data, the greater the distortion in the distribution of expected results using the independent historical return method. This point is confirmed in Table III which compares the actual/expected distribution of results for the 912 monthly small company stock returns, which have historically shown a higher degree of variability than large company stock returns.⁴ As shown in Table III, the actual/expected difference for the independent historical return method rises to 328, and a visual inspection of the results again clearly shows the overstated variability. Table III also shows that using a conditional probability method reduces this actual/expected difference to 101.2, and provides a much more consistent distribution of results.

Not only is the distribution of monthly results much more consistent with the actually observed returns, the expected 76-year wealth becomes the actually observed 76-year historical value of \$7,860; the geometric mean becomes the actually observed 12.5%, and the expected arithmetic mean becomes 15%, a number which again compares favorably with an actual arithmetic mean which can vary from 12.6% to more than 17% depending upon how the data is used to determine the arithmetic mean.⁵

Comparable results for the independent historical return method using monthly data to determine parameters are expected wealth of \$193,611; expected geometric mean of 17.4%, and an expected arithmetic mean of 22.6%, if the expected wealth is actually achieved.

V. STOCK OPTION PRICING

The last two sections showed the accuracy of a forecast model with lognormal parameter selection based on the conditional probability method, as compared with a forecast model with the lognormal parameters based on the independent historical return method. However, it took several years in order to accumulate the data necessary to demonstrate the problems with the independent historical return method by showing that the expected distribution of monthly returns is a poor match with the actual history.

Such an accumulation of historical data would not be needed with a stock option pricing model. If the model is based on the independent historical return method for parameter selection, and the conditional probability method is the correct approach, then the stock option pricing model would immediately provide a poor match with actual market-based stock option quotations.

The best-known stock option pricing model which is based on the independent historical return method of parameter selection is the Black-Scholes Stock Option Pricing Model. The Black-Scholes Option Pricing Model is an analytical tool that is widely used to estimate the theoretical value of a stock option.⁶ The model uses five key determinants: stock price, strike price, volatility, time of expiration, and short-term (risk free) interest rate. The model was published in the Journal of Political Economy in 1973 and has been considered as a major break through in the field of option pricing. This model is, in fact, experiencing some accuracy problems. In particular, Mark Rubinstein in a 1994 Presidential Address to the American Financial Association noted that: "...Black-Scholes has become increasingly unreliable over time in the very markets where one would expect it to be most accurate . . ."⁷; and David S. Bates documented many problems with the Black-Scholes model in his 2002 paper, Empirical Option Pricing: A Retrospection.⁸

However, when the parameter selection process is changed from the independent historical return method to the conditional probability method, and another technical problem in the Black-Scholes model is corrected, the result is a Conditional Probability Method (“CPM”) Stock Option Pricing Model which replicates actual market prices quite well.

Since this section is relatively long it is helpful to divide the presentation into six separate subsections, each of which is briefly summarized as follows:

- Section V(i): Theoretical Option Pricing Model develops the basic theoretical framework for the adjusted option-pricing model.
- Section V(ii): Black-Scholes Comparison highlights the differences between the Black-Scholes model and the theoretical framework presented in Section V(i).
- Section V(iii): Risk-free Rate of Return presents the role that the risk-free rate of return plays in option pricing, and how it affects option pricing models.
- Section V(iv): CPM Stock Option Pricing Model sets out the details of the new stock option pricing model.
- Section V(v): Empirical Data Comparison compares the predicted option prices with empirical results.
- Section V(vi): Stock Option Pricing Commentary presents some additional commentary on the relevant parameters in a stock option pricing model.

V(i). THEORETICAL OPTION PRICING MODEL

This section will be used to develop theoretical option pricing models for both a stock call option and a stock put option.⁹ A stock call option is a contract that gives the holder the right to purchase a share of stock at a pre-established price. The common use of the stock call option in an employee compensation package is for the employer to grant the employee the right to purchase shares of stock in the company at a fixed price, called the strike price. For example, if the stock is currently selling at \$80 per share, the employee might be granted the right to purchase shares at a fixed price of \$90 per share (the strike price). The incentive for the employee is to help boost company performance, which will then be reflected in the company’s stock price. Should the share price exceed the strike price threshold (for example, it rises to \$105), the employee could then exercise his or her option making an immediate gain (\$15 in the case of the example) on each stock call option exercised.

A stock put option is a contract that gives the holder the right to sell a share of stock at a pre-established price. Stock put options can be used to help insure an investor against a market loss.

The function of an option-pricing model is to put an estimated value on the contract holder’s right to purchase or sell shares at a fixed price. This right has a value even if shares are currently selling for less than the strike price (in the case of a call option) or more than the strike price (in the case of a put option), because of the potential that share value may change.

The value of a stock option depends on the distribution of anticipated changes in the value of the stock. For example, for a call option if there is a relatively high probability that the stock price will grow to exceed the strike price, or even a small, but significant, probability that the stock price will grow to exceed the strike price by a large amount, the option clearly has more value than if there is a high probability the share price will always remain below the strike price, or if in those instances where the share price is expected to exceed the strike price, the excess is relatively small.

The first step in building a theoretical stock call option pricing model is to calculate for any given time horizon, t , the expected distribution of return relatives. Assuming that returns are lognormally distributed, and letting x be a given return relative, the lognormal density function which describes this distribution is as follows:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}\sqrt{t}} \frac{1}{x} e^{-\frac{(\ln x - \mu t)^2}{2\sigma^2 t}} \quad (17)$$

This formula is similar to the lognormal density function presented in equation (1). The only difference is that it adds a new parameter, t , to reflect the stock option time horizon.

The second step is to calculate the expected future share value given an initial share value of S . But this is simply Sx , or the initial share value multiplied by the appropriate rate of growth in stock value.

The third step will be to calculate the potential gain from a call option, assuming that the gain in the stock price of x is actually achieved. If the strike price is K then the theoretical gain from the stock call option is $Sx - K$, for any return, x , that is greater than K/S .

The fourth step is to calculate the theoretical gain from the stock call option. This is the probability weighted summation of all possible $Sx - K$ gains, and is given by the formula:

$$\int_{K/S}^{\infty} (Sx - K) \frac{1}{\sigma\sqrt{2\pi}\sqrt{t}} \frac{1}{x} e^{-\frac{(\ln x - \mu t)^2}{2\sigma^2 t}} dx \quad (18)$$

Finally, since the above expression represents the expected gain at the end of the option expiration period, the theoretical stock call option price is then just the discounted value of the above expected gain using an appropriate discount rate, r . Thus, the theoretical stock call option price is given by the following equation:

$$OP_c = e^{-rt} \int_{K/S}^{\infty} (Sx - K) \frac{1}{\sigma\sqrt{2\pi}\sqrt{t}} \frac{1}{x} e^{-\frac{(\ln x - \mu t)^2}{2\sigma^2 t}} dx \quad (19)$$

While formula (19) seems relatively complex, an interesting development occurs if the assumed annual discount rate, r , is set equal to the expected annual rate of return for the underlying security, which is just the expected value for the distribution of returns described by the density function, $f(x)$, in equation (17). The equality between these two rates is expressed as follows:

$$r = \mu + \frac{\sigma^2}{2} \quad (20)$$

In this case, the theoretical stock call option price as described by formula (19) becomes:

$$OP_c = SN(d_1) - Ke^{-rt}N(d_2) \quad (21)$$

Where N is the cumulative standard normal distribution, and

$$d_1 = \frac{\ln(S / K) + (r + \frac{\sigma^2}{2})t}{\sigma\sqrt{t}} \quad (22)$$

$$d_2 = d_1 - \sigma\sqrt{t} \quad (23)$$

This formula is, of course, the same formula that appears in Black-Scholes Option Pricing Model.¹⁰ The theoretical development presented in this section is somewhat different than the presentation offered by Black and Scholes in their 1973 paper. It emphasizes the importance of equation (20) in the valuation of stock options using the Black-Scholes methodology. Given that Black and Scholes derived the option pricing formula by a different approach, the importance of equation (20) appears not to have been addressed.

There is an exactly parallel development for the theoretical price of a stock put option. In this case equation (20) plays the same important role, and the theoretical stock put option price is given by:

$$OP_p = SN(d_1) - Ke^{-rt}N(d_2) - (S - Ke^{-rt}) \quad (24)$$

where d_1 , and d_2 are defined exactly as described above.

A comparison between the formulas developed above and the traditional Black-Scholes Option Pricing Model is provided in the next section.

V(ii). BLACK-SCHOLES COMPARISON

Section V(i) developed a theoretical option pricing model which generates the exact Black-Scholes Option Pricing Model formula. It uses all of the same inputs as the Black-Scholes model, but with one key difference – the discount rate, r , and the lognormal distribution parameters, μ and σ , must be selected with due regard for equation (20) in order for the formula to be algebraically correct. In the Black-Scholes model the discount rate, r , is set at the risk free rate, which would usually be somewhat lower than the expected rate of return on the underlying security,

$$\mu + \frac{\sigma^2}{2} .$$

A key advantage to the theoretical development presented in Section V(i) is that it clearly emphasizes that in the Black-Scholes formula there is a relationship between the discount parameter, r , and the volatility measure σ . If each of these parameters is calculated independently, this relationship could be lost causing the option valuation results to be skewed.

It is noted here that the risk-free rate of return is a critical element of the final option pricing model presented in this paper, however the discussion of its role will be presented in Section V(iii).

This still leaves open the important question of calculating the lognormal parameters, μ and σ , from historical data. The Black-Scholes Option Pricing Model is based on the independent historical return method and uses σ as described by equation (8) as its volatility parameter. The option pricing model presented in this paper determines a new volatility parameter, v , which is based on the conditional probability method. Under the conditional probability method, the parameters μ and σ^2 as determined by equations (7) and (8) are reduced proportionately until the expected return for the lognormal distribution is equal to the expected long-term geometric mean return for the underlying security. Letting g be the long-term expected geometric mean, under the conditional probability method, g is simultaneously equal to the expected discount rate, r , and the expected lognormal parameter μ , as calculated by equation (7) for a very large sampling of historical data. These relationships dictate the proportional reduction so that the new volatility parameter, v , is described as follows:

$$v = \sqrt{\frac{2g\sigma^2}{2g + \sigma^2}} \quad (25)$$

where σ is as defined in equation (8). Determining the volatility parameter, v , in this fashion and using it in lieu of σ in the basic option pricing formulas always assures that the relationship between the expected return and expected volatility as described by equation (20) holds true.

It is still necessary to highlight the role that the risk free rate of return does play in option pricing. That topic is addressed in the next section of this paper.

V(iii). RISK-FREE RATE OF RETURN

Before describing a final stock option pricing model, it is necessary to discuss the importance of the risk-free interest rate, and the put-call parity theorem.¹¹ The put-call parity theorem states that for a given strike price the difference between an actual call price and an actual put price must equal the difference between the current share price of the underlying security, and the strike price discounted at the risk-free rate of return. If this equation does not hold, the situation leads to an arbitrage opportunity providing a guaranteed return for the investor.

In the theoretical development presented so far, the relationship between call and put prices fails the put-call parity theorem; thus, one or both of the above theoretical option prices need to be adjusted so that this theorem holds true.

The put-call parity theorem, and the use of the risk-free rate of return, only addresses the difference between a given put price and a given call price, not their specific values. For example, the put-call parity theorem may dictate that the difference between two given prices be \$5, but does not specify whether the actual prices are \$20 and \$25, \$24 and \$29, or some other pair of prices whose difference is \$5. Thus, the theoretical development presented in Sections V(i) and V(ii) is still a critical element of the final option pricing model presented in this paper.

By examining empirical data it is evident that for strike levels somewhat below the current stock price, the put option branch of the theoretical formula presented in Section V(i) tends to match actual put option prices. In this situation when the theoretical call option prices are adjusted to comply with the put-call parity theorem, they also match actual call option prices.

However, at strike levels somewhat above the current stock price, the reverse is true. At the higher strike levels, the call option branch of the theoretical formula presented in Section V(i) tends to match actual call option prices. And once again when the theoretical put option prices are adjusted to comply with the put-call parity theorem, they too match actual put option prices.

These results are exactly what would be expected. For relatively low strike levels, the theoretical put option price approaches zero. Thus, it is not possible to adjust the theoretical put option prices, and any adjustments must be made to the call option formula. For relatively high strike levels, the reverse is true, since at high strike levels the call option prices approach zero.

Accordingly, the model presented in the next section includes adjustments in the theoretical option prices to reflect the put-call parity theorem. For purposes of the illustrations presented in Section V(v), the theoretical call option formula will reflect the entire put-call parity adjustment if $Ke^{-gt}/S \leq 1.01$. If $Ke^{-gt}/S \geq 1.10$ then the theoretical put option formula will reflect the entire put-call parity adjustment. Pro rata adjustment will be provided to each theoretical formula for values of Ke^{-gt}/S between 1.01 and 1.10.

Clearly, more research is called for to help identify when and how these put-call parity adjustments should be made. However, in the mean time, practitioners concerned with FAS 123(R) compliance may wish to apply the above processes to develop specific models based on data from traded options, and then use the model to estimate the value of options that are not traded. With the above concepts in mind, the next section of this paper provides the specific details of a new CPM Stock Option Pricing Model.

V(iv). CPM STOCK OPTION PRICING MODEL

This section presents the details of a new stock option pricing model. The model looks somewhat similar to the Black-Scholes model, but has three key differences:

1. the risk-free rate of return parameter, r , has been replaced with an expected long-term rate of return parameter, g , so that the importance of equation (20) in the algebraic development of the formula is recognized;
2. the volatility parameter, σ , is replaced with a new parameter v , as described by equation (25), to reflect the conditional probability method of parameter determination; and,
3. an adjustment factor has been added to each theoretical formula, so that the requirements of the put-call parity theorem are met.

The model is parameter driven and relies on eight key variables, which are described as follows:

S = current stock price.

K = strike price.

t = time remaining until expiration expressed in years.

r = current continuously compounded risk-free interest rate.

g = current continuously compounded expected long-term rate of return for the underlying security. This would normally range from around 10% ($g = .09531$) for stocks of larger companies to around 11.5% ($g = .10885$) for stocks of smaller companies.

σ = volatility as measured by the standard deviation of the logarithms of historical returns, expressed as an annual percentage. This is the same volatility measure that is currently used in Black-Scholes calculations.

A = the strike level at which the put-call parity adjustment begins to be transferred from the call option formula to the put option formula. A is expressed as a percentage of the strike level at which the unadjusted theoretical call option price equals the unadjusted theoretical put option price.

B = the strike level at which the put-call parity adjustment is fully transferred from the call option formula to the put option formula. B is expressed as a percentage of the strike level at which the unadjusted theoretical call option price equals the unadjusted theoretical put option price.

Other notation that is used in the model is described as follows:

\ln = the natural logarithm function

N = the cumulative normal distribution

e = the base of the natural logarithm ($e = 2.7183$)

The specific formula for the estimated put option price is:

$$OP_p = SN(d_1) - Ke^{-gt}N(d_2) - (S - Ke^{-gt}) + K(e^{-rt} - e^{-gt})(1-w) \quad (26)$$

The specific formula for the estimated call option price is:

$$OP_c = SN(d_1) - Ke^{-gt}N(d_2) - K(e^{-rt} - e^{-gt})(w) \quad (27)$$

Where: $d_1 = \frac{\ln(S/K) + (g + \frac{v^2}{2})t}{v\sqrt{t}} \quad (28)$

$$d_2 = d_1 - v\sqrt{t} \quad (29)$$

$$v = \sqrt{\frac{2g\sigma^2}{2g + \sigma^2}} \quad (30)$$

$w =$ weight parameter. This parameter expresses the percentage weight given to the call option formula in order to reflect the put-call parity adjustment factor. The parameter $w=1.00$ if $Ke^{-gt}/S \leq A$, $w=0.00$ if $Ke^{-gt}/S \geq B$, and w is calculated by linear interpolation for values of Ke^{-gt}/S between A and B .

Empirical data illustrations are completed in the next section of this paper.

V(v). EMPIRICAL DATA COMPARISON

The stock option pricing model presented in this paper has been entirely theoretical. As with any theoretical model, it is important to compare predicted output using reasonable assumptions with actual observed results.

This paper compares actual market prices with prices predicted by both the current Black-Scholes model and the CPM model described above in Section V(iv). The first comparison is based on S&P 500 Index Options. The data for the comparison shown below comes from the Business Section of The Washington Post for August 16, 2005, which lists prices for 13 different 66-day S&P 500 Index options. The relevant assumptions used for the comparison are:

Estimated Volatility (σ):	16%
Estimated Long-Term Rate of Return:	10%
Estimated Current Risk-Free Rate of Return:	3%
Adjustment Parameter A:	1.01
Adjustment Parameter B:	1.10

Empirical Data Comparison
(S&P 500 Index: 1,234)

Call Option Comparison

<u>Strike Price</u>	<u>Washington Post</u>	<u>New Formula</u>	<u>Current Black-Scholes</u>	<u>Difference between Actual Prices and Formula</u>	
				<u>New Formula</u>	<u>Current Black-Scholes</u>
1,200	\$52.60	\$51.12	\$56.94	\$ 1.48	\$ 4.34
1,225	34.20	33.90	41.52	.30	7.32
1,245	21.00	21.94	31.32	.94	10.32
1,250	19.00	19.39	29.00	.39	10.00
1,260	14.70	14.45	24.76	.25	10.06
1,275	9.00	8.84	19.41	.16	10.41
1,280	7.50	7.56	18.02	.06	10.52
1,285	6.30	6.66	16.22	.36	9.92

Put Option Comparison

<u>Strike Price</u>	<u>Washington Post</u>	<u>New Formula</u>	<u>Current Black-Scholes</u>	<u>Difference between Actual Prices and Formula</u>	
				<u>New Formula</u>	<u>Current Black-Scholes</u>
1,200	\$12.00	\$10.74	\$16.57	\$ 1.26	\$ 4.57
1,215	18.20	14.86	21.81	3.34	3.61
1,225	18.40	18.40	26.02	.00	7.62
1,245	27.50	26.33	35.71	1.17	8.21
1,250	29.00	28.75	38.37	.25	9.37

Average Difference: \$.77 \$8.17

The second comparison is based on The Wall Street Journal Listed Option Quotations for November 4, 2005 which lists eight different 77-day options for Intel Corporation. The relevant assumptions based on historical data and reasonable expectations are as follows:

Estimated Volatility (σ):	40%
Estimated Long-Term Rate of Return:	10%
Estimated Current Risk-Free Rate of Return:	3.5%
Adjustment Parameter A:	1.01
Adjustment Parameter B:	1.10

Empirical Data Comparison
(Intel Share Price: \$23.99)

Call Option Comparison

<u>Strike Price</u>	<u>Wall Street Journal</u>	<u>New Formula</u>	<u>Current Black-Scholes</u>	<u>Difference between Actual Prices and Formula</u>	
				<u>New Formula</u>	<u>Current Black-Scholes</u>
\$20.00	\$4.20	\$4.22	\$4.45	\$.02	\$.25
22.50	2.05	2.16	2.64	.11	.59
25.00	.68	.78	1.40	.10	.72
27.50	.15	.36	.66	.21	.51

Put Option Comparison

<u>Strike Price</u>	<u>Wall Street Journal</u>	<u>New Formula</u>	<u>Current Black-Scholes</u>	<u>Difference between Actual Prices and Formula</u>	
				<u>New Formula</u>	<u>Current Black-Scholes</u>
\$20.00	\$.10	\$.08	\$.32	\$.02	\$.22
22.50	.40	.50	.99	.10	.59
25.00	1.50	1.61	2.23	.11	.73
27.50	3.70	3.68	3.97	.02	.27

Average Difference: \$.09 \$.49

V(vi). STOCK OPTION PRICING COMMENTARY

When the Black-Scholes Model was published in 1973:

- a. The volatility parameter was based on the assumption that historical investment return observations constitute an independently determined set of data.
- b. The discount rate was set at the risk-free rate of return because of the potential for arbitrage.

This paper has demonstrated that there are problems with each of the above processes. The reason for the change to the conditional probability method for volatility parameter determination is completely explained in Sections II, III and IV. The fact that the CPM Model replicates actual market option prices just confirms that the conditional probability method of parameter determination is appropriate.

The source of the discount rate problem can be traced back to the original Black and Scholes 1973 paper. Unfortunately, the presentation included in the 1973 paper had an unstated assumption that there is one, and only one, investment return related parameter. Thus, the presentation precluded the possibility of a stock option pricing model which reflects the dynamic interaction of two different investment return related parameters. The theoretical development presented in this paper emphasizes the need for both the risk-free rate of return and the expected long-term rate of return on the underlying security to be included in a stock option pricing model.

Any calculation problems caused by the above original Black-Scholes processes tended to offset one another in actual computations. Thus, computation problems with the formula were not immediately detected. This paper has demonstrated that the basic formula itself is still theoretically correct and has modified the original formula into a new model. With this model practitioners can derive a simple tool for estimating stock option prices. While

additional research may help to clarify the exact put-call parity adjustment process, the new model helps practitioners deal with the immediate concerns of complying with FAS 123(R).

The new model also provides analysts with a tool to easily determine investor expectations based on actual market activity in traded options. Once parameters are selected such that the model output matches actual listed option quotations, the parameter g provides the market's expectation of the rate of return for the underlying security, based on actual market quotes for options.

VI. INVESTOR IMPACT

The ideas presented in this paper could have a significant impact on actual investor behavior. Should the conditional probability method of analyzing historical data be adopted, it would directly impact the calculation of risk premia, or additional expected investment return for taking additional investment risk, as measured by the volatility of returns. In this section, two common premia are recalculated using the conditional probability methodology to show the magnitude of the change advocated.

The two risk premia that are recalculated are: (1) the equity premium, or the additional return expected for investing in large company stock instead of treasury bills; and (2) the small stock premium, or additional return expected for investing in small capitalization stocks instead of large capitalization stocks.

The premia comparison is as follows¹²:

	Independent Historical Return Method	Conditional Probability Method
Equity Risk Premium	8.6%	6.6%
Small Stock Premium	3.3%	1.6%

The conditional probability method calculations reflect the actually observed historical differences. In other words, over the past 76 years the actual small company stock returns have exceeded the actual large company stock returns by an average of 1.6% per year, and the actual large company stock returns have exceeded the actual treasury bill returns by an average of 6.6% per year. Yet, the methodology employed by the independent historical return method of parameter selection produces an expected small stock premia of 3.3%, and an expected equity risk premia of 8.6%, and these results are communicated to investors considering the purchase of large or small company stocks.¹³ It is also important to note that because of the process used in the independent historical return method, the expected future excess return will always exceed the actual historical result, no matter how long the history is.

While no one can predict the future, and some investors might be completely comfortable with the independent historical return method, perhaps other investors would be interested in at least knowing this difference between actual and projected results exists, and in these cases the difference between predicted results and actual historical results should be communicated to the investors.

No such communication is necessary under the conditional probability method. The expected small stock premia is determined in such a fashion that it equals the actually observed historical premia. The conditional probability approach provides a method of demonstrating risk/reward for potential investors without using a model which always produces results which are significantly larger than those that are actually observed.

A second area of impact could be on the calculation of beta, the capital asset pricing model (CAPM) well-known and commonly used estimate of market risk.¹⁴ Attached as Table IV is an illustration which shows how different methods of calculation can produce significantly different beta estimates. The illustration shows a 10-year history of both the market proxy and a specific investment, each of which yield the same wealth over the 10-year period. Using the complete history of quarterly returns yields a beta of .388, whereas using annual snapshots of the history yields a beta of 1.548, a significantly different result.

While no calculation philosophy will completely eliminate the differences between any two separate calculations based on the same data, the recognition that historical returns have a conditional nature can significantly reduce the magnitude of the difference. The recognition of the conditional nature requires the analyst to separate observed

volatility into two components, one which adds to wealth and one which does not. Beta can then be calculated such that it reflects only the first of these components.

VII. SUMMARY

American workers are relying more and more on defined contribution approaches toward meeting their retirement objectives. In addition, strong consideration is being given to altering Social Security so that it, too, has a large defined contribution element. In this environment, workers need complete and accurate information which show the potential impact of alternative savings and investment strategies, so that they can adequately prepare for the future.

However, the lognormal model with parameters determined under the independent historical return method ignores the fact that wealth creation is essentially a continuous process.¹⁵ In this setting, when one takes a limited number of observations and treats them as independently determined rates of return, the resulting analytical tool may be used to make forecasts. However, the forecasts vary widely depending upon how the data is used, and in all cases the long-term expected results are larger than the historical values upon which the model was based.

The independent historical return method generates an expected return equal to the “arithmetic mean of the historical results.” Unfortunately, the “arithmetic mean of historical results” is not unique for any given history, and may vary significantly depending upon how the history is expressed. For example, Table IV in the appendix shows two different investment histories, one for a market proxy and one for a specific security. But as shown below, these histories have exactly opposite “arithmetic means” depending upon the time period for measuring return.

Table IV Arithmetic Means (Annualized)

Measurement Period	Market Proxy	Specific Security
Quarter	15.0%	12.0%
Year	12.0%	15.0%

When the conditional nature of observed periodic returns is considered, the expected value of the annual rate of return equals the geometric mean of the entire data set when that data set consists of a single multi-year observation. The recognition of this fact, and the change to conditional probability methodologies, could impact investor decisions since the risk element of the decision does not change, but the reward element decreases.

In addition, when the conditional nature of observed periodic returns is considered, it is possible to develop a stock option pricing model which consistently replicates actual market quotations. Such a model will be a big asset to companies concerned with FAS 123(R) compliance. It will also be helpful for analysts who wish to gain as much information as possible from implied expected results based on actual market equations for stock options.

Table I

Comparison of Actual Monthly Returns with Expected Monthly Returns for Large Company Stocks

Independent Historical Return Method – (Parameters Based on Monthly Data)

Rate Range	Expected Number of Results	Actual Number of Results	Difference
.9389 or less	91.2	67	24.2
.9389 to .9623	91.2	64	27.2
.9623 to .9794	91.2	76	15.2
.9794 to .9943	91.2	95	3.8
.9943 to 1.0085	91.2	118	26.8
1.0085 to 1.0229	91.2	128	36.8
1.0229 to 1.0385	91.2	117	25.8
1.0385 to 1.0570	91.2	126	34.8
1.0570 to 1.0832	91.2	77	14.2
1.0832 or more	91.2	44	47.2

Total difference between expected and actual: 256.0

Table II
 Comparison of Actual Monthly Returns with Expected Monthly Returns for Large Company Stocks
 Conditional Probability Method

Rate Range	Expected Number of Results	Actual Number of Results	Difference
.9478 or less	88.0	83	5.0
.9478 to .9678	88.7	67	21.7
.9678 to .9825	89.4	76	13.4
.9825 to .9953	90.3	86	4.3
.9953 to 1.0074	90.8	104	13.2
1.0074 to 1.0196	91.0	108	17.0
1.0196 to 1.0328	91.4	101	9.6
1.0328 to 1.0485	92.4	116	23.6
1.0485 to 1.0707	93.9	94	0.1
1.0707 or more	96.1	77	19.1

Total difference between actual and expected: 127.0

Table III

Comparison of Actual Monthly Returns with Expected Monthly Returns for Small Company Stocks

Rate Range	Actual Number of Results	Independent Historical Return Method Expected Results	Independent Historical Return Method Difference	Conditional Probability Method Expected Results	Conditional Probability Method Difference
.9401 or less	116	182.4	66.4	108.7	7.3
.9401 to .9657	71	91.2	20.2	101.4	30.4
.9657 to .9882	106	91.2	14.8	118.1	12.1
.9882 to 1.0098	130	91.2	38.8	126.8	3.2
1.0098 to 1.0318	149	91.2	57.8	126.5	22.5
1.0318 to 1.0558	136	91.2	44.8	118.4	17.6
1.0558 to 1.0847	99	91.2	7.8	102.8	3.8
1.0847 or more	105	182.4	77.4	109.3	4.3

Total difference between expected and actual:

328.0

101.2

The lognormal parameters for the independent historical return method were determined using traditional methods and monthly data. The results are: $\mu = .00974$ and $\sigma = .08495$. The degree of distortion for this method is reduced somewhat if the parameters are determined as monthly equivalent results based on annual data, but the same pattern of overstated variability emerges. The conditional probability method parameters used for this illustration are $\mu = .00803$ and $\sigma = .06011$

Table IV

Market Proxy and Specific Security Returns for Beta Illustration

Year End	Quarter End	Market Proxy Value	Specific Security Value
0	4	1.000	1.000
1	1	1.200	1.142
	2	1.560	1.304
	3	1.560	1.489
	4	1.500	1.700
2	1	1.210	1.561
	2	1.210	1.434
	3	1.375	1.317
	4	1.210	1.210
3	1	1.513	1.339
	2	1.815	1.482
	3	1.906	1.640
	4	1.573	1.815
4	1	1.494	1.720
	2	1.465	1.630
	3	1.737	1.545
	4	1.464	1.464
5	1	1.830	1.684
	2	1.922	1.869
	3	1.729	1.992
	4	1.757	2.050
6	1	1.932	2.234
	2	1.836	2.257
	3	2.279	2.077
	4	1.772	1.772
7	1	2.392	2.215
	2	1.914	2.365
	3	2.067	2.525
	4	2.038	2.303
8	1	2.384	2.395
	2	2.193	2.352
	3	2.245	2.311
	4	2.144	2.144
9	1	2.723	2.679
	2	3.077	2.943
	3	3.138	3.232
	4	2.930	3.118
10	1	3.194	2.977
	2	3.162	2.843
	3	3.525	2.716
	4	2.594	2.594

ENDNOTES

1. See FASB (2004).
2. See Ibbotson Associates (2002).
3. See Joss (2005) for development of the 10.8% to 13.8% range.
4. See Ibbotson Associates (2002) for the small company data.
5. See Joss (2005) for the development of the 12.6% to 17.0% range.
6. See Bodie (2001).
7. See Rubinstein (1994).
8. See Bates (2002).
9. The model presented in this section assumes that the stock does not pay dividends. To adjust the model presented in Section V(iv) to reflect expected dividend payments, let the parameter g be the expected long-term rate of return less the portion of this return attributable to the expected dividend rate.
10. See Black (1973) or Bodie (2001).
11. See Bodie (2001).
12. These results are based on the data in Ibbotson Associates (2002). Using more recent data would alter the specific results, but not the relationship between the results as measured by the two different parameter determination methods.
13. See Shoven (2004).
14. See Graham (2001).
15. See Joss (2004) and Joss (2005).

REFERENCES

- Bates, David S., Empirical Option Pricing: A Retrospection, (University of Iowa and the National Bureau of Economic Research), 2002.
- Black, Fischer and Myron Scholes, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, Vol. 81, No. 3, May/June 1973, 637-654, 1973.
- Bodie, Zvi, Alex Kane and Alan J. Marcus, Essentials of Investments, Third Edition (McGraw-Hill/Irwin, Boston), 2001.
- Financial Accounting Standards Board (FASB), FAS 123(R) Share-Based Payment, (FASB, Norwalk, CT), 2004.
- Graham, J.R. and C.R. Harvey, The Theory and Practice of Corporate Finance: Evidence from the Field, *Journal of Financial Economics*, 60, 187-243, 2001.
- Ibbotson Associates, Stocks, Bonds, Bills and Inflation 2002 Yearbook (Ibbotson Associates, Chicago), 2002.
- Joss, Richard, Correcting the Overstatement in Investment Forecasts, *Journal of Financial and Economic Practice*, Fall 2004, 47-53, 2004.
- The Tendency of the Arithmetic Mean to Overstate Expected Returns, *Journal of Financial and Economic Practice*, Fall 2005, 72-87, 2005.
- Rubinstein, Mark, Presidential Address: Implied Binomial Trees, *The Journal of Finance*, Vol. 49, No. 3, July 1994, 771-818, 1994.
- Shoven, John B., The Trade-Off Between Risk and Reward, *On Investing*, Winter 2004, 16-17, 2004.

Richard R. Joss is a retired actuary living in Bethesda, Maryland. He enjoys working on mathematical problems, riding his bike, hiking and wood carving. He can be reached on the internet at rrjoss@starpower.net.