Accession Capital Corp

R-SQUARED & T-STATISTICS AS A MEASURE OF STOCK BETA ACCURACY

By Richard R. Conn CMA, MBA, CPA

EXECUTIVE SUMMARY

- This paper is a continuation of "Sensitivity of a Stock Beta …"
- R-Squared and T-Statistics only measure the linearity between a Market Index and the movements of a subject stock price
- It is impossible for any regression to separate the systematic and non-systematic elements
- Improved R-Squared and T-Stats may be occurring by random chance, even when the distance between the predicted and true beta is widening

This paper is a continuation of the earlier Accession "Sensitivity of a Stock Beta to Non-Systematic Price Impacts". A full appreciation of the intricacies of this subject matter would require that the reader first review that earlier work upon which this investigation builds. However, a synopsis of that work will be provided here in order to provide a foundation for this study.

REVIEW OF "Sensitivity of a Stock Beta to Non-Systematic Price Impacts"

In that earlier work[1] the general question was raised 'How much do non-systematic price shocks to a stock price distort the quality of the CAPM Beta (β) regressed from the combined data?' In order to address that question a hypothetical equity was created that moved in lock-step with the S&P/TSE Index at a predetermined β of 1.3. Not surprisingly, when that equity was subject to linear regression against the Market Index, the slope of the β line was found to be exactly 1.3 with an R^2 of precisely 1.0 Then, in order to test the sensitivity to non-systematic price distortions, a random element was introduced to the month-end test equity outcomes. Initially, a random amount of +/- 10%

[1] Available on the Accession Capital Corp website at http://www.accessioncap.com

variation was introduced to equity price[2]. Now, the equity data with the non-systematic 'noise' added was again regressed against the Market Index and the resultant observed β was compared to the "pure" 1.3 starting point.

In general, it was found that even small amounts of non-systematic noise caused considerable β prediction errors in the regression of single equities.

PURPOSE OF THIS INVESTIGATION

The question raised in the previous work but only partially addressed there was 'Is there some quantifiable means by which a regressed β might be tested for dependability?' Empirically it was shown that a collection of regressed equities will provide a weighted average β significantly closer to the 'true' un-distorted β of the group because the non-systematic error terms tend to cancel each other out[3].

Moreover, the Coefficient of Determination (or R-Squared statistic) was shown in the previous paper to provide some means of adjudging how well the regression had fitted the Beta-line to the distorted data. The R-Squared results were shown to be an imprecise indicator of β quality. That is to say, there were instances where samples with a higher R^2 produced a regressed weighted-average β farther away from the "true" hypothetical β of 1.3. This is because R^2 is strictly a measure of the linearity between the Market Index and actual Stock Price – not a means of testing how much of the stock price movements are the result of systematic vs. non-systematic price impacts. The corollary here is that, if a given regression produces a very low R^2 we would be well advised to disregard the results entirely on the premise that the regression was not able to determine a strong

[2] For example, if the Market Index gained 1.0% in any given month, the equity price appreciation could be expected to randomly fall anywhere in the [(1+(1.0% x 1.3)) x (0.9 to 1.1) - 1] range.
[3] The important caveat here, however, is that all the equities in the group sampled were hypothetical and unrealistically had identical "true" betas. This is unlikely to ever occur in the real world no matter how homogeneous the firms may be in any given industry. The complexity of blending a weighted average of differing Betas is considered in the adjunct paper: "MEASURING THE ERROR OF ESTIMATION IN GROUPED STOCK PRICE BETAS".

linear relationship between Market Index movements and Stock Price fluctuations[4]. We cannot, however, conclude that a regression producing a high R^2 is predicting a more precise β estimate than one with a lower R^2.

So, the purpose of this paper is to consider if the often cited "T-Statistic", perhaps in harmony with the R^2 measure can provide greater insight into the fidelity of the regressed β. Since even small inaccuracies in the estimation of the Beta can have a truly profound impact on the assessment of a firm's cost of capital, or a project's discount rate, any means of improving the estimation process will come as a real boon to the valuator's art.

It is likely that, for most working finance professionals, some years have passed since last having been in an introductory Statistics class. Therefore, a brief refresher on the basic construct of a T Statistic and the probabilities under a Normal (Gaussian) Distribution Curve, as well as the concepts surrounding the Null and Alternative Hypothesis are presented in APPENDIX ONE.

THE T-STATISTIC

The purpose of the T-Statistic, as applied to regression analysis, is to test the Null Hypothesis that there is no linear relationship between the two variables, the Market Index and the subject Stock Price. If this were to be true, the actual Beta would be expected to be zero. So the T-Statistic tests the likelihood of arriving at the regressed Beta value given the assumption that the true Beta is zero. It is the actual distance between the regressed Beta and zero, divided by the standard error of the regression that attests to the statistical probability that there is no linearity between the Market and the Stock Price movements. Intuitively, one must believe that, the farther the absolute value of the regressed Beta is from zero, and the smaller the value of the standard error of the regression, the less likely that the Null Hypothesis can be accepted.

[4] This point was tacitly made in the first paper when, as the general level of monthly random potential noise was added to the data sample (from 2% to 5 to 10 to finally 15% plus or minus "true"), the simple average of R-Squared declined from 0.90 to 0.67 to 0.28 to 0.20 respectively. See Table B in that work.

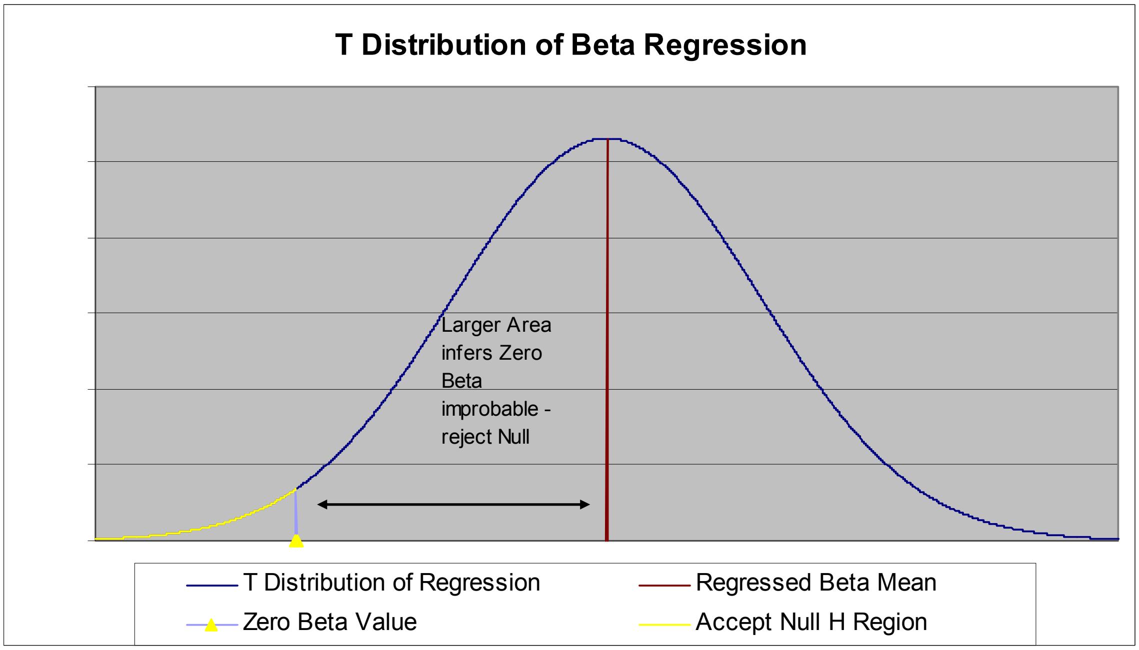

However, when the Accept Null Hypothesis region[5] is proportionately much larger:

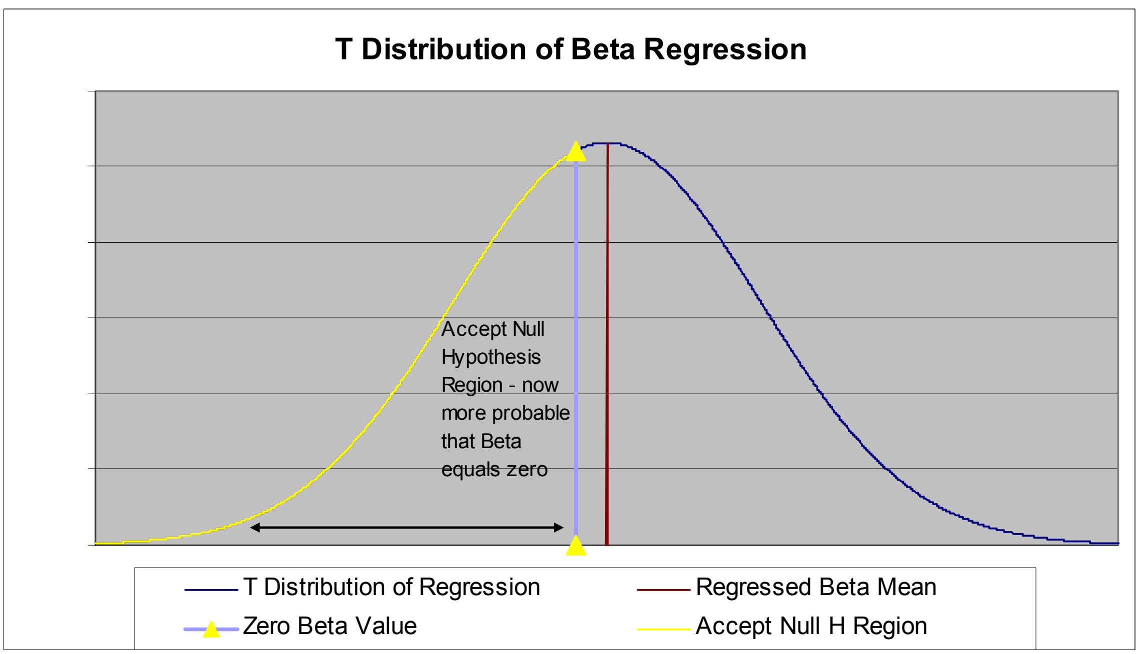

[5] Since we are testing the hypothesis that the actual Beta equals zero, a one-tailed test is required. We have assumed that the regressed Beta is positive, and therefore the Accept Null Hypothesis region is on the left side of the graphics presented.

It becomes much more probable that the actual Beta is zero and the Null Hypothesis should therefore be accepted.

Note that the T-Statistic in no way assesses the precision of the regressed Beta. The only qualitative information that can be inferred from the T-Statistic (and the related P-Value) is how probable it may be that the regression was completed in the absence of any linear relationship between the variables.

Using exactly the same data as presented in Table A of the "Sensitivity of Stock Beta …" paper, we now include the T-Statistic and associated P-Value:

TABLE A – SIMULATED 60 MONTH PRICE HISTORY w +/- 10% Random Noise

Instance	Observed β	R-Squared	T-Statistic	P-Value
TEST A	0.964	0.188	3.659	0.2740%
TEST B	1.344	0.337	5.435	0.0001%
TEST C	1.248	0.299	4.973	0.0003%
TEST D	1.613	0.400	6.217	0.0000%
TEST E	**1.290**	0.282	4.777	0.0006%
TEST F	0.896	0.204	3.859	0.0144%
TEST G	1.137	0.245	4.338	0.0029%
TEST H	**1.290**	0.306	5.055	0.0002%
TEST I	0.840	0.206	3.878	0.0136%
TEST J	1.273	0.331	5.351	0.0001%
Simple Average	1.190	0.280	4.754	0.0306%
Minimum Observed	0.840	0.188	3.659	0.0000%
Maximum Observed	1.613	0.400	6.217	0.2740%
Range	0.773	0.212	2.558	0.2740%

Again, using the Empirical Rule as a general guide, we know that almost 99% of all observations would be expected to fall within plus-or-minus three standard deviations from the mean. So, it is not surprising the P-Values are so very small, given that the lowest T-Statistic is almost 3.7 standard errors. P-Value, in this case, represents the "Accept Null Hypothesis" area in the tail of the distribution. Accordingly, a higher T-Statistic necessitates a lower P-Value. The application of the T-Statistic to this data confirms that there is a very high linear relationship in all the regression tests presented. In fact, as an overall simple average, the P-Value is telling us there is only a $3/100^{th}$ of 1% probability that these T-Statistics could have been observed if, in fact, the “true” Beta were really zero. Practically speaking, however, the T-Statistic provides little other qualitative insight into the precision of the Betas observed.

In fact, because we know that all of the aforementioned ten regression tests were developed from a “true” Beta of 1.3, we can see that ‘Test D’ demonstrates both the highest R-Squared as well as the highest T-Statistic, but ranks as the forth *worst* predictor of true beta, being 0.313 away from 1.3. Tests E & H are tied for the position of best Beta predictor (each only 0.01 away from 1.3) and yet the associated R-Squared and T-Statistics are approximately reporting in about the midpoint on the overall range of R-Squared and T-Statistics.

Note, as well, that the T-Statistic is ordinal with the values of R-Squared. That is, the instance with the lowest value of R-Squared (Test A) also reports the lowest T-Statistic. This correspondence in ranking continues in sequence through all 10 instances where Test D has both the highest R-Squared as well as the highest T-Statistic.

CONCLUSIONS

The T-Statistic when applied to regressed stock betas (where the “true” systematic only Beta is never known) can only test the linearity between the Market Index and subject stock price movements. Even when there is a high degree of linearity between the two variables, neither the R-Squared value nor the T-Statistic can be relied upon to determine the precision of the Beta estimate. A higher R-Squared / T-Statistic might be indicating greater fidelity with the unknown underlying “true” Beta, or these higher values may be happening purely by random chance given that the non-systematic stock price distortions have coincidentally happened to improve the linearity between the two regression variables.

APPENDIX ONE – THE T STATISTIC

Review: Normal Probability Function

The Normal Probability Density Function (i.e. the Gaussian Distribution) is a symmetrical curve defined by $\int_{Guass}(x)$ is drawn on an 'x/y' Cartesian grid system[6]. The area under the total curve is said to represent the probability that a given outcome will occur – with the most likely or *expected* outcome represented by the midpoint of the distribution where the height of the curve is greatest. If, for example, a large population of stocks, known to next be paying a dividend between $0.00 and $2.00 is identified, and a randomly selected sub-group is selected from that population, then the distribution of the average dividend per share could be:

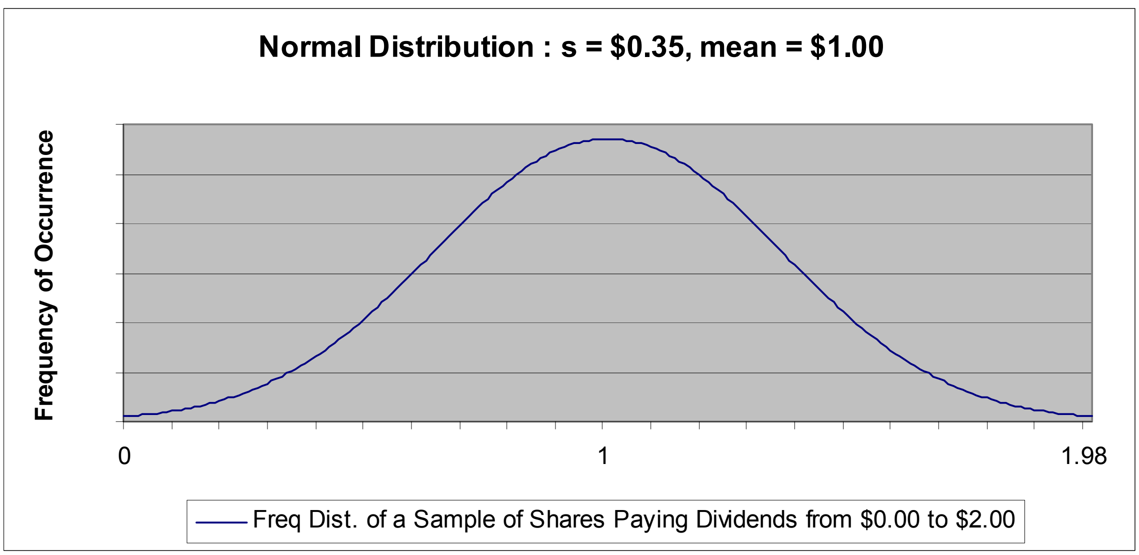

The "expected" average dividend per share would be $1.00. In the distribution presented, there would be an equal probability of selecting a subgroup of stocks paying less than $1.00 than a group paying more than $1.00. This is because the area under the curve up to $1.00 is exactly 50% of the total area. Moreover, it should be remembered that there is a relationship between the area under the curve (and, therefore, the probability of occurrence) and the standard deviation. The empirical rule, for example, says that approximately 68% of all outcomes would be expected to occur within plus-or-minus one

[6] The formula for the Normal Probability Density Curve is $\int_{Guass}(x) = (1/(\sigma\sqrt{2\pi}))e^{-(x-\mu)^2/(2\sigma^2)}$ where σ is the standard deviation of x and μ is the mean of x.

standard deviation from the population mean. In our dividend example, then, we would expect the majority of all samples selected to display an average dividend per share of between $0.65 to $1.35 ($1.00 +/- $0.35). Deductive reasoning tells us that there is only a 16% probability [(100% - 68%)/2] of randomly selecting a group of stocks that, on average, paid less than $0.65/share. And, equally, there is only a 16% probability of selecting a sample of shares that paid more than $1.35.

CONSTRUCT OF THE BETA SAMPLE

In the case of the per-stock averages, the Population was the known group of equities that next planned to pay dividends in the $0.00 $2.00 range and standard deviation could be precisely determined. The inference we were attempting to make was how closely any given sample average would reflect the true known $1.00 mean of the population.

In the case of attempting to make inferences about the quality of a given stock beta regression, there are many differences. The data is fixed in time, historic and immutable. We are examining the relative change in an equity price given a known change in the Market Index. Theoretically the Population could be said to be entire realm of divisible moments during the time period under study. We could, for example, measure the Beta of a stock price on month-end values, weekly, daily hourly … etc. Contrary to accepted norms of many precision standards, increasing the sampling frequency of a stock beta is not expected to improve the accuracy of the observations.

One very important reason this is so is because the impact of non-systematic price distortions are cumulative and non-correcting. Take, for example, a stock where the "true" beta is believed to be 1.0 (and the relative stock price movements therefore should mirror the percentage change in the index) and sample points are taken weekly. At the start of the measurement period the Index is 100 and the stock price is $1.00. At the end of the first week the Index finishes up at 101 and, correspondingly, the stock price is up

1.0% at $1.01. On the next day of trading, the Index does not move, however, negative non-systematic news causes the stock price to fall back to $1.00. Subsequently, at the end of that week the Index moves up to 102 and the stock price moves in harmony with the market to $1.01 (approximately another 1% increase in both the market and the stock). On a week-over-week basis, the Index has climbed from 101 to 102, but the stock price has remained static at $1.01. When included in the regression, all other things being equal, the derived Beta will now be less than 1.0 because there has been an occasion where the Index value moved, but the stock price seemingly did not. Increasing the sampling frequency to daily would not necessarily improve the beta precision because then the countervailing distortion of the day when the Index remained unchanged, but the stock price dropped as a result of non-systematic news would then be included in the sample.

Unlike the average-dividend experiment, therefore, we cannot simply increase the sample size in order to increase the probability that the resultant mean of the sample will be a true reflection of the population mean. Generally, a beta regressed from a historic period of 60 periods of month-end data is held to be most reflective of the firm's actual beta[7]. These 60 months are regressed against corresponding month-end movements in the Market Index and the result is one single Beta statistic that represents the outcome for that sample of data. The hypothetical (but immeasurable) Population is the conceptual Beta that would have resulted from a sample where the error term caused by non-systematic stock price changes could magically be stripped out from the pre-regression data.

[7] Implicit in this practice is the assumption that the firm has not undergone some major reorganization, acquisition or divestiture during those sixty months that would cause the inherent risk-level of the firm to change.

THE T-STATISTIC

Conceptually the T-Statistic works under the same principals as the simple Empirical Rule example provided above. That is, an estimate of the Population Standard Deviation (symbolized σ) is derived and then any given observation can be stated in terms of how many standard deviations that observation is away from the mean. The area under the curve (and therefore, probability of occurrence) at any given factor of standard deviations can easily be calculated because, in a normal distribution, area is a function of Mean and Standard Deviation.

The important distinction in our application of the T-Statistic is that neither the Population Mean nor σ is actually known. As a result, regressed Beta (the sample mean) is used as a proxy for the unknown Population Mean (the "True" Beta), and a point estimate of the standard error in Beta Deviation is used as a substitute for the unknown Population Standard Deviation[8].

The formula is:

$$T_{obs} = (\text{в} - \beta) / \left[\left(\sqrt{\left(\Sigma (y_i - \hat{y}_i)^2 / (n - 2) \right)} \right) / \sqrt{\left(\Sigma (x_i - \bar{x})^2 \right.} \right]$$

Where:

T_{obs} is the observed "T" value at that sample Beta
в is the regressed Beta value
β is the Population Beta (unknown)
y_i is the ith observation of stock price delta
$\hat{y}_i$ is the ith y value on the regressed Beta line [will be equal to $\alpha + \text{в}_{xi}$ where α is the regression intercept and x_i is the value of x at the ith observation]
n is the number of observations in the regression (e.g. 60 months)
x_i is the ith observation of the Market Index delta
$\bar{x}$ is the mean of all the x observations

and note that all of

[8] While not germane to our purposes here, it can be noted that when the degrees of freedom (related to the size of the sample) are small, the actual shape of the T-Distribution differs from the Normal Distribution. However, once the sample size becomes large enough, the T and Normal Distribution share that same distinctive Bell-Shaped curve.

$$[(\sqrt{(\Sigma(y_i - \hat{y}_i)^2 / (n-2))}) / \sqrt{(\Sigma(x_i - \bar{x})^2}]$$

is really just a means of approximating the standard error of the Beta Deviation

If, for example, we knew that the "true" β of a population was 1.3, and our regression estimated a в of 1.6, the T Statistic would tell us how many estimated standard deviations 0.3 Beta (1.6 – 1.3) was away from the mean and the area under the T Distribution curve to the right of this demarcation point would indicate how probable obtaining such a regression outcome would be.

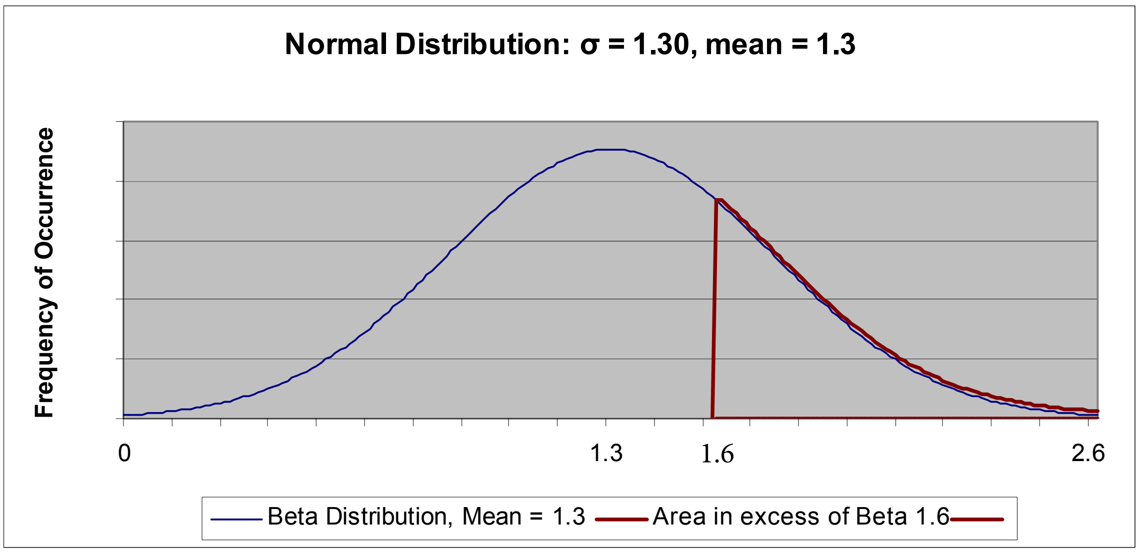

We do not, however, know the Population β – that is, in fact, the inference we are attempting to make from the regression sample. So the best inference the T Statistic can provide us in the circumstances relates to how likely it is that there is a linear relationship between the movement of the subject Stock Price and the Market Index.

NULL AND ALTERNATE HYPOTHESIS

The logic goes as follows: If there was NO linear relationship between Market Index and Stock Price, the β of the Stock Price ("y" dependent variable) relative to the movement of the Market Index ("x" independent variable) would be zero. That is, the change in the Stock Price given a shift in the Market Index would be entirely sporadic and random.

The Null Hypothesis becomes: β = 0, The Alternative Hypothesis, therefore, becomes β <> 0, in which case there must be some linearity in the relationship between the subject Stock and the Market Index.

Using our previous example of a 1.6 ʙ sample regression (which in this case we know is not the true Beta), the same process is repeated where (1.6 – 0) beta is converted into the number of standard deviations away from the sample mean of 1.6. The probability of being this many standard deviations away from the estimated mean is determined by examining the area under the T Distribution Curve at that point and then making an inference about the validity of the Null Hypothesis.

If, for example, it were found that a distance of 1.6 ʙ was calculated to be 5 standard deviations away from the estimated mean, we would have to conclude that it would be highly improbable that a regressed sample of 1.6 beta could have been obtained under the condition that the actual β = 0 (Remember, under the Empirical Rule, 99% of all outcomes can be expected to be found within +/- 3 standard deviations from the mean … so a distance of 5 would be highly unlikely). Under such conditions we would have to reject the Null Hypothesis and therefore conclude that there was, indeed, some linear covariance between the Market and the Stock Price.

Note that there is nothing particularly qualitative about the T Statistic. The purpose of the comparison is merely to assess whether it is likely that a linear connection exists between the x and y variables. We would probably have much more faith in a ʙ that was calculated to be 5 standard deviations away from the point where β = 0 than one which was only one standard deviation from the β = 0 delineation. However, in neither case is the T Statistic testing the precision of the ʙ regressed, only the likelihood or the remoteness of obtaining such an outcome in the event that there was no linear relationship between the two variables.

P-VALUE

P-Value is simply the area under the curve beyond a given standard deviation (or standard error in the case of a T-Statistic). Specifically, P-Value is the proportion of that area to the total area under the curve – and this, therefore, represents the probability of occurrence for any standard error to be observed in this region. In the Beta Distribution example depicted above, the area under the red line represents the probability of observing a Beta regression in excess of 1.6.