Historic Stock Price Volatility and its Implications to Business Valuators

By Richard R. Conn CMA, MBA, CPA, ABV

EXECUTIVE SUMMARY

- Stock price volatility is an important indicator of asset risk, but historic volatility profiles provide qualitative, rather than quantitative information.
- Graphing the volatility history of the target firm simultaneously with that of its competitors and the Market Index can provide unique insight into the risk/comparative advantages of the target.
- Other things being equal, a firm with a higher volatility can be said to be 'riskier', but the differential in risk is not necessarily proportional with the difference in volatility.
- Historic asset volatility includes both systematic and non-systematic price impacts. The Market Index volatility, in contrast, will show very little non-systematic risk.
- Volatility is a simple metric and does not convey the 'drift' or geometric rate of growth/decay in the asset.
- There is no proportional relationship between an asset's volatility and its beta.

INTRODUCTION

Intuitively, we expect there is a direct relationship between risk and expected stock price volatility[1]. Risk-free debt instruments have zero volatility[2] and predicting their value at any future point in time is relatively easy. At the other extreme we know that assets with extremely high volatility have, by definition, a wide range of market prices over a comparatively short period of time. 'Risk', then, can be thought of as the degree of certainty we might have in estimating a future expected value for any given asset. It stands to reason that we will have a much higher probability of accurately estimating the future value of an asset that is expected to incur only a 10% annual rate of volatility compared with one with 100%.

Business valuators are often required to offer their opinion as to the cost of capital on publicly traded assets. Further, the cost of capital on a publicly traded asset can sometimes serve as a proxy for a

[1] In this paper we will use the terms 'volatility' and 'standard deviation' interchangeably. Typically, when financial professionals and the press reference the volatility of a given stock price, they mean the standard deviation of that equity's total rate of return measured at some appropriate frequency (hourly, daily, weekly, monthly) and stated at an annualized rate. More insight into the measurement process is provided in the "Mechanics of Volatility" section.
[2] Ignoring interest rate risk / reinvestment risk for simplicity.

similar asset that is not publicly traded. Therefore, any added insight that the historic volatility of a publicly traded asset can offer into the riskiness of that asset will enhance the depth of analysis the valuator can bring to the question at hand.

BACKGROUND – VOLATILITY IS STOCHASTIC

Contrary to the traditional Black Scholes assumption of a constant volatility for the underlying security, financial engineers have pointed out for many decades now that the implied volatility of virtually all exchange traded options are anything but constant. Implied Volatility is solved for by holding the observed Black Scholes inputs of Current Spot Price; Strike Price; Risk-Free Rate; and Dividend Yield as independent constants and determining the volatility rate that would be required to equate all these fixed variables to the current actual option price. In the Black Scholes equation, the volatility assumption is monotonic, meaning that there will be one and only one unique volatility rate that will equate all the independent variables to the observed market option price.

Typically what is most often observed in the implied volatility market data is a "smile". That is, even with the same time to maturity, any series of actual Puts or Calls, differing only by strike price, will show higher implied volatility at the extreme ends of the strike price. For example, the following series of Puts on Royal Bank of Canada stock (which ranged in strike price from $42 to $70), all of which had approximately 7 weeks to maturity at the time this data was captured showed an implied volatility profile that ranged from a low 22.6% to a high of 51.6%. Either the Market is significantly mispricing these derivatives, or the notion of a constant volatility just has no bearing in reality.

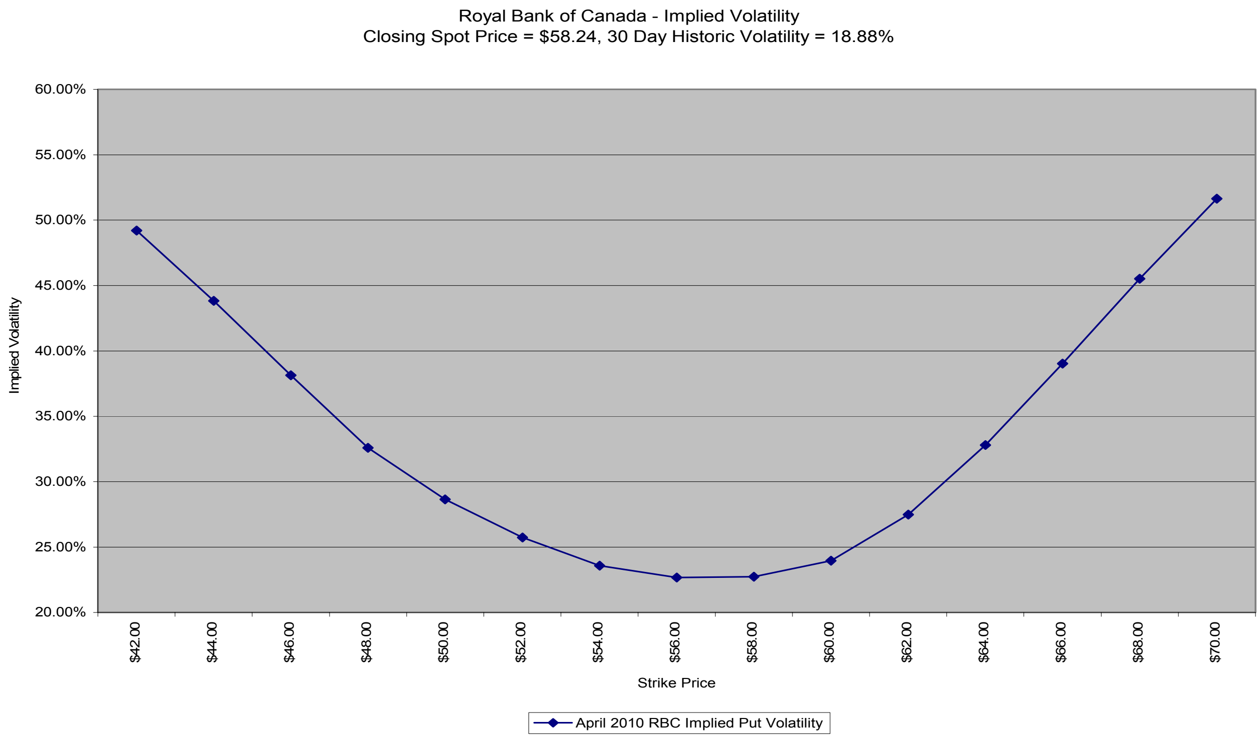

Further, if the implied volatility is plotted for a series of differing strike prices on a number of options with sequentially increasing maturity dates, a 'volatility surface' can be developed that will provide a complex snapshot of how the market perceives the underlying security's volatility to change over time *and* strike price. The important point for our purposes here, however, is to note that almost never is the implied volatility surface flat or constant as the traditional Black Scholes model would have us believe.

One of the reasons proposed for non-constant future volatility expectations is that volatility is believed to be a stochastic (random) function (see, for example, [1]). Clearly, for hedge fund managers and arbitragers the existence of a non-constant volatility has enormous implications that must be factored into every investment decision. Is the same true for business valuators?

If the Market believes that future rates of volatility for any given stock is increasingly random (i.e. are a stochastic function of both the length of time period being projected as well as the potential price differential away from the current spot price), then, do historic stock price volatilities convey *any* meaningful information to the business valuator?

Further, is it, for example, reasonable to assume that, all other things being equal, a security that has consistently shown an annual stock price volatility of 40% is more risky than one that historically is only at 20%? And, if this is reasonable, is the first twice as risky as the second? Moreover, is there a proportional relationship between the size of a security's historic volatility and its beta?

THE MECHANICS OF VOLATLITY

The Standard Deviation of any sample is simply the square root of that sample's Variance. Variance, in turn, is simply the averaged sum of the squared distances of every observation from the mean of the sample:

$$\text{Variance} = [\sum (y_i - \hat{y})^2] / (n - 1) \text{, where } \hat{y} \text{ is the sample mean}$$

So it is important to note that already we are talking about an averaging process (i.e. dividing by n – 1 observations). Therefore, the more observations we have in our sample of stock price movements, the less probable it will be that a few rouge outliers will have a significant impact upon our results.

Remember too that we are only concerned with *relative* stock price movements here – the rate of investment return over the period of observation. While it is possible to measure relative stock price changes in terms of the actual linear change over a period[3] as in:

$$\text{Delta \% Rate of Change} = (\text{Period2 Price} / \text{Period1 Price}) - 1 = \text{Observation } y_i$$

It is much more common to measure the natural logarithm of the period change:

$$\text{Delta \% Rate of Change} = \text{Ln}(\text{Period 2 Price} / \text{Period 1 Price}) = \text{Observation } y_i$$

[3] The issue of measurement frequency will be discussed in APPENDIX A. Relative stock price changes here would include total returns including the impact of dividends.

SOME EMPERICAL EVIDENCE

The Toronto Stock Exchange divides its listed companies into seven subsectors: Clean Tech; Diversified; Energy; Income Trusts; Life Sciences; Mining; and Technology firms. The 252 day rolling volatility for each of the largest[4] capitalized firms in each of these sectors has been plotted against the volatility of the TSX Stock Index as a whole:

[4] Market equity capitalization as at Dec. 31, 2009 are reported by the TSX was used as the selection criteria. In the case of the Clean Tech sectors, the 2nd largest firm was selected as a result of historic data quality issues with the largest firm.

252 Day Rolling Volatilties - TSX Sector Leaders

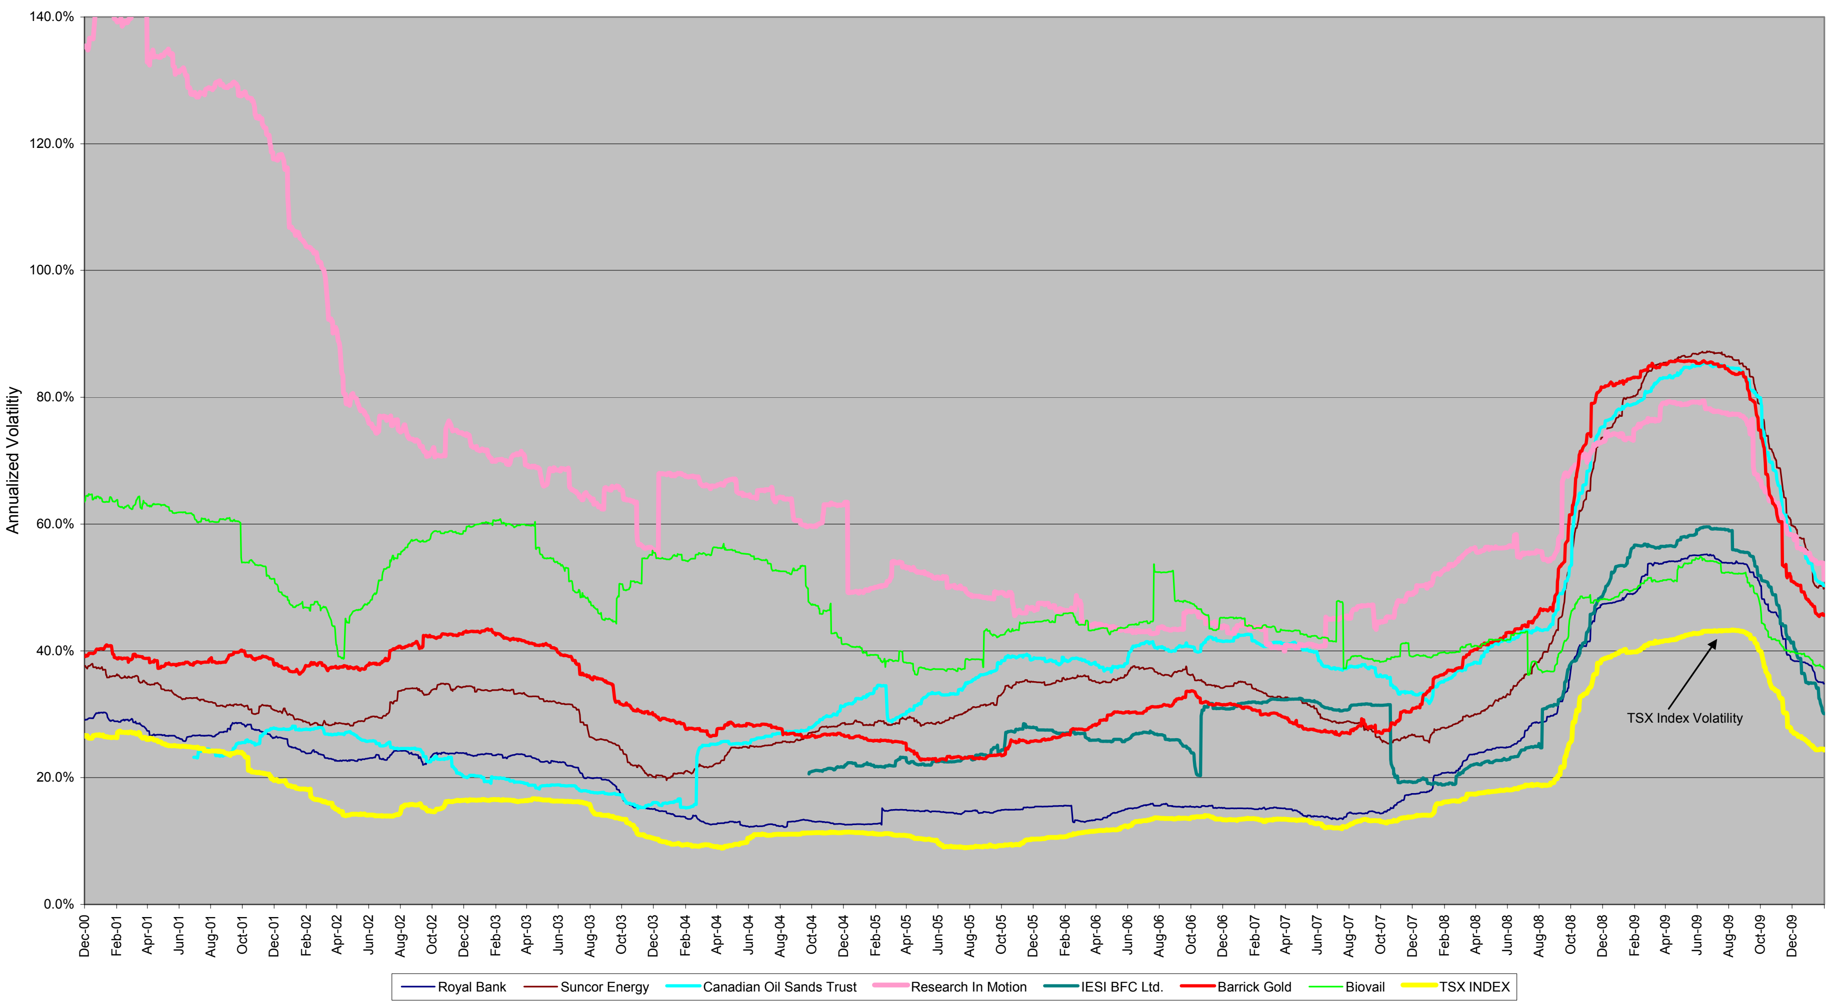

Some salient observations should become immediately apparent:

- The TSX Index consistently displays the lowest volatility (this is in spite of the fact that at least one of these equities, the Royal Bank, can be expected to have a lower beta than the Market Index – the relationship between beta and volatility will be discussed later).
- With the exception Research In Motion and Biovail, the rest of the equities can be generally said to 'trend' along with the TSX Index
- The 2008 financial economic crisis is plainly evident in the marked increase in post September 2008 volatilities.
- There is clear evidence of 'jump diffusion' in five of the seven equities. Four of these raise interesting cost of capital issues and will be discussed further.
- **On the whole, with few exceptions, the progression of annual average standard deviations remains quite predictable and centered around a relatively narrow range.**

This last point is of particular interest to business valuators as it specifically speaks to the issue of implied future volatilities. "The empirical findings indicate that the conditional variance, log-variance, and standard deviation of stock market returns are pulled back to some long-run average level over time.[5]" This indicates that stock volatilities are mean-reverting over a longer term. For the arbitrager, attempting to quantify volatility over the next day, week, month or quarter, the process is stochastic and fraught with uncertainty. In contrast, the business valuator can probably rely upon the average of the past five years of historic standard deviations as a reasonable indicator of future *long-term* expectations. The evidence indicates that the mean-reverting tendencies of stock price volatility will eventually smooth out any short-term randomness.

The exception to this general stance is when the equity in question is clearly in a period of transition. Research In Motion (RIM) is a prime example. By the start of the new millennium RIM was just beginning to enjoy its meteoric stock price appreciation as its revolutionary Blackberry product gained increasing acceptance in the marketplace. In the intervening years RIM's products have been adopted as indispensible business tools and the firm's financial future

[5] See [2] at pg. 19

has become much more secure[6]. It would be highly unlikely that RIM's pre-2005 volatility could be presumed to be indicative of future performance. While we have not provided other comparative examples, that earlier volatility profile is expected to be characteristic of innovative post-development high-tech firms that have incurred extremely high rates of growth during the first-to-market product adoption phase. Eventually, however, technology substitutes will enter the market and slow the overall rate of growth, as will market saturation.

Obviously, as well, the volatility of any individual security, or the Market Index from September 2008 through February 2010 should be considered circumspect as representative of the future long-term mean. But this period does provide us with a useful range of parameters to estimate future conditions in the event that another global economic shock is to be endured.

If we examine the volatility of the volatilities (i.e. the standard deviation of the 252 day rolling standard deviations – but excluding the days after August 2008 in order to eliminate the recession distortions), a better appreciation of the long-term predictability of volatility can be had. Column C of TABLE 1 reports one standard deviation of the 252 day equity volatilities.

TABLE 1

	JAN 2000 - AUG 2008 SAMPLE		A	B	C	C / A
SECTOR	NAME	SYMBOL	Average Volatility	# of Observations	σ of 252 day Rolling σ	σ as a % of the Average Volatility
Mining	Barrick Gold Corp	ABX	34.5%	2,177	7.2%	21.0%
Clean Tech	IESI BFC Ltd.	BIN	25.7%	981	4.2%	16.3%
Life Science	Biovail	BVF	49.0%	2,177	7.9%	16.2%
Income Trst	Canadian Oil Sands Trust	COS.UN	30.7%	1,795	8.3%	27.2%
Technology	Research In Motion	RIM	73.5%	2,177	31.4%	42.7%
Integrated	Royal Bank of Canada	RY	19.8%	2,177	5.7%	28.8%
Energy	Suncor	SU	31.2%	2,177	4.4%	14.0%

Note that, in the case of the Royal Bank, one standard deviation equates to 5.4%, which, by application of the empirical rule, means that we can be 95% sure that the actual expected volatility of RY will lie within +/- 2 standard deviations of the sample average of 19.8%. In

[6] As much, that is, as any firm's future can be assured. It was not long ago that using the terms "General Motors" and "Bankruptcy" in the same sentence would have amounted to corporate heresy.

other words, we can be 95% certain that actual RY volatility lies between 8.4% and 31.2% over the long-term. Contrast this with the relative uncertainty in RY *implied volatility* over the seven weeks ranging from 22.6% to 51.6% (as indicated from the Put option discussed on pg. 2).

Predicting the expected value of any stochastic process is much easier in the long-run compared with the short-run if that process is mean-reverting.

Not all equities 'trend' with the market volatility. Smaller firms (who generally can be expected to have a higher proportion of non-systematic risk to total risk than compared with their larger counterparts) will often show historic volatilities relatively unrelated with the market[7]. For example:

[7] Some care needs to be exercised here, however, because one reason why a small firm volatility does not track with the market might be due to the fact that it is too thinly traded to properly reflect its actual price volatility.

Accession Capital Corp

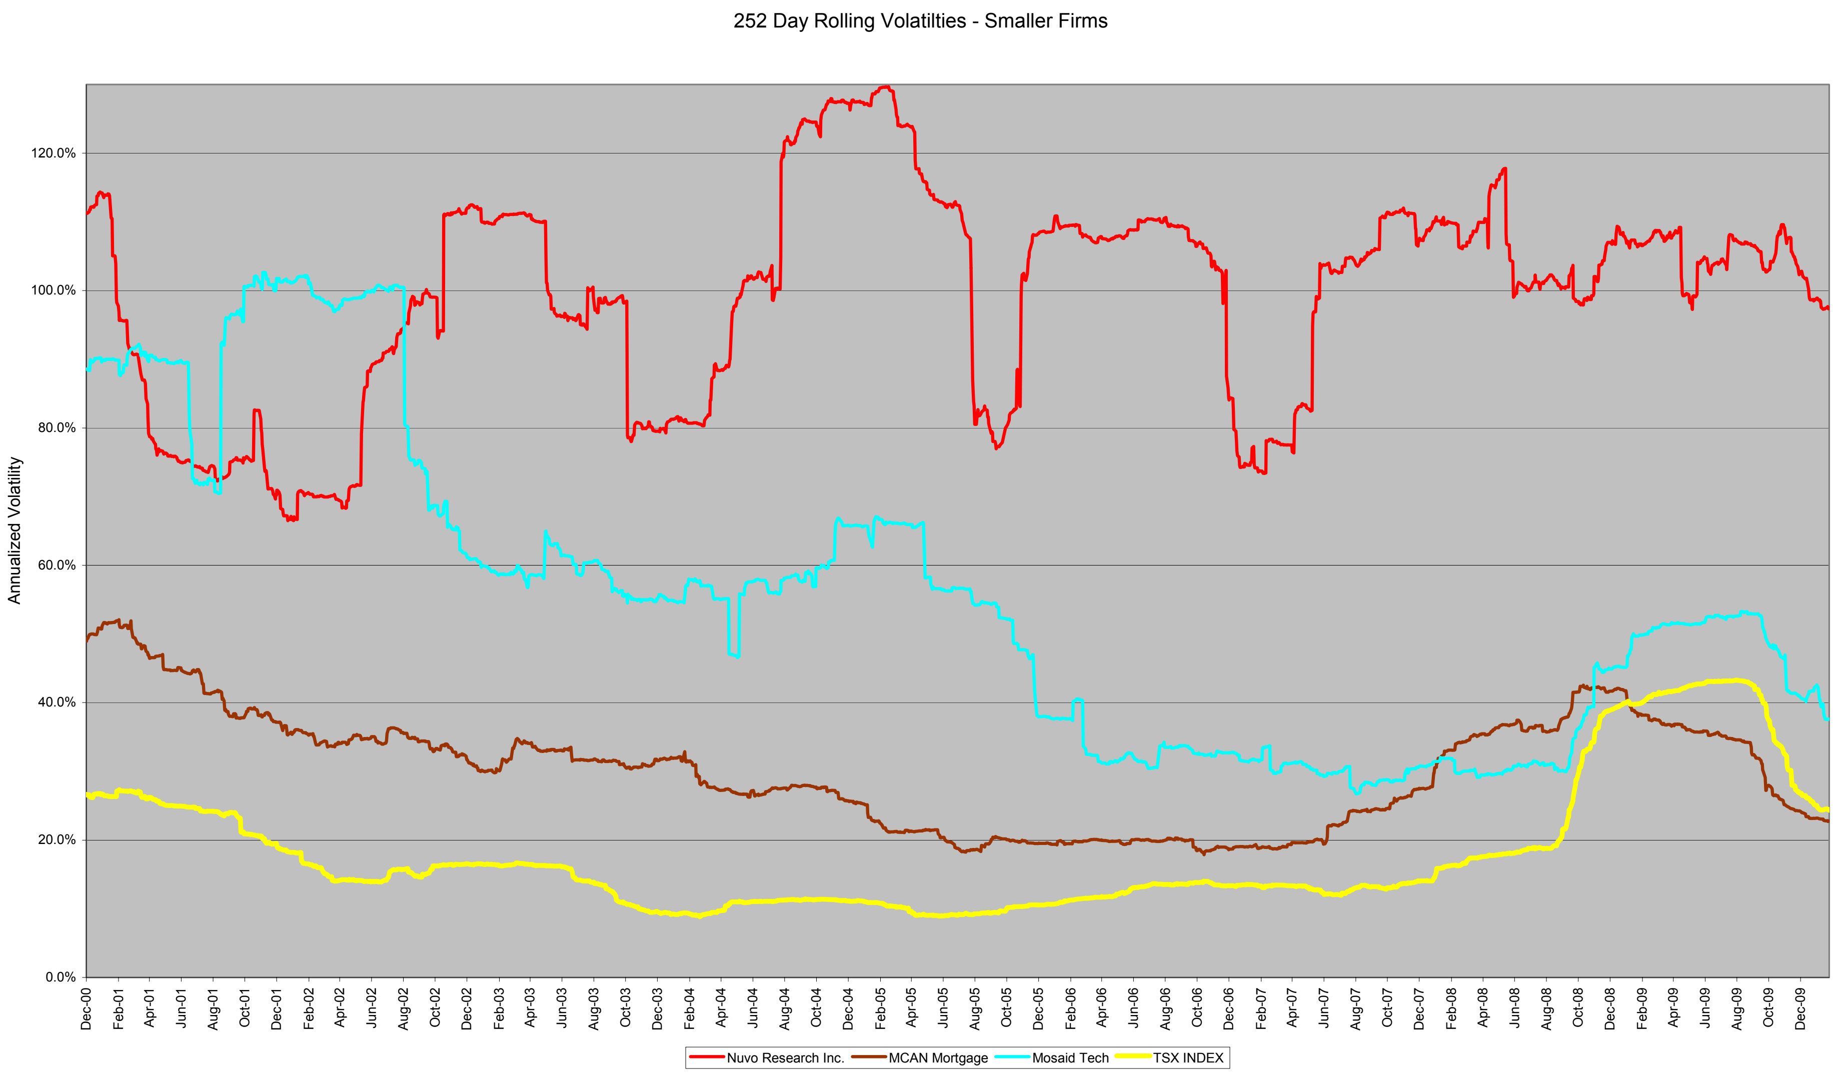

Clearly, MCAN Mortgage, Mosaid Techologie and Nuvo Research Inc. are much less correlated with the Market volatility than, for example, the Royal Bank is. In fact, the correlation coefficient between Nuvo and the Market Index is less than 0.0044 compared with that of Royal Bank's 0.975. While it would be difficult to quantify exactly how this fact impinges upon the relative cost of capital of the two firms (other than, of course, to conclude that Nuvo must face much more non-systematic risk than Royal Bank does), it is a point worth considering as a qualitative factor in the determination of Nuvo's cost of capital.

MCAN Mortgage is an interesting anomaly for two reasons: it shows only a modest increase in volatility between June 2007 and October 2008, a time when the US securitized real estate market was self-imploding; and, at the height of the US financial crisis MCAN volatility actually descends below that of the TSX as a whole. The relative trend of MCAN's declining volatility compared with the financial industry as a whole is something that should be given qualitative consideration in any study of MCAN's costs of capital[8]. The fact that MCAN's volatility has remained so controlled probably indicates a superior business strategy and/or a well-executed hedging position. Compared with Quest Capital Corp., a peer firm, that also invests only in Canadian mortgages and is relatively of the same market capitalization as MCAN:

[8] MCAN invests primarily in high quality Canadian only residential and commercial mortgages. They source funds from short to mid Term Deposits that are CDIC insured, thereby avoiding the necessity to offer the small debt investor (< $100K) any form of collateral. The common shareholder consistently earns dividends that are taxed as interest, but at yields well in excess of 10%. All these factors seemed to have combined to make the MCAN common share price very stable in spite of the tumultuous financial markets. Note that Royal Bank, Canada's largest chartered bank, incurred substantial increases in volatility post August 2008 (as did all the chartered banks – the data simply has not been graphed herein), whereas MCAN's volatility began descending at this point.

Accession Capital Corp

252 Day Rolling Volatilties - 2 Smaller Mortgage Firms

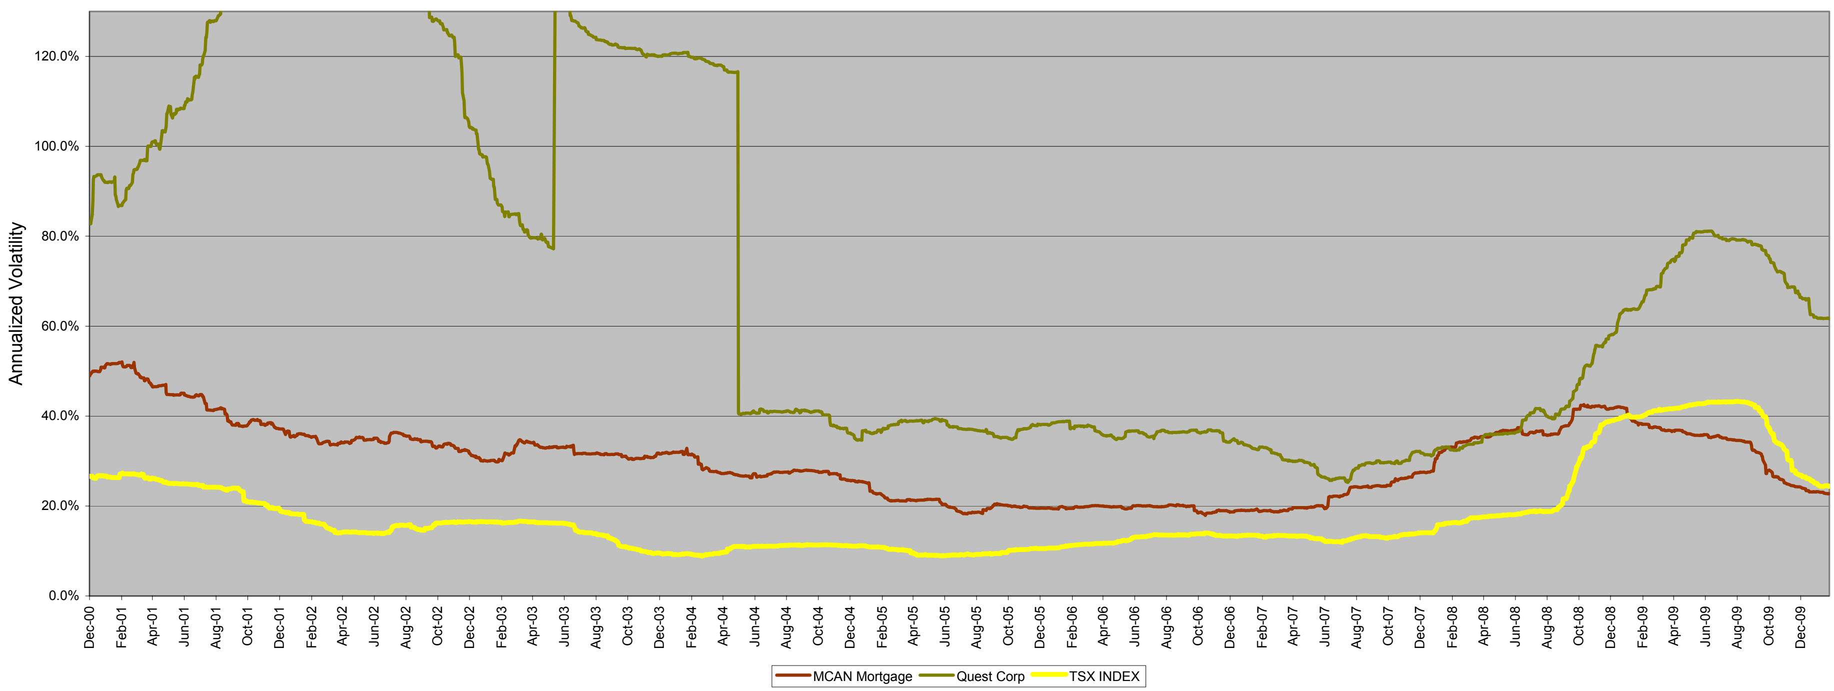

We see that MCAN has virtually always enjoyed a lower volatility than Quest which would, ceteris paribus, imply that it probably also enjoys a lower cost of capital (On July 4, 2003 Quest announced a merger/acquisition transaction that substantially increased its 10 year average volatility [see Jump Diffusion section]. Even so, Quest has proved to have the higher volatility in almost all the periods subsequent to the July 2003 anomaly). The intuition that MCAN would have a lower cost of equity capital than Quest appears to be borne out by the regressed betas. MCAN shows a beta of 0.53 whereas Quest shows a beta of 1.43[9]:

TABLE 2

Comparison of Beta @ Period of Best R-Squared

	MCAN Mortgage	Quest Capital Corp
Observed Beta	0.5348	1.4314
R-Squared	0.1242	0.3836
Standard Error of Y Estimate	0.0408	0.0603
36 Month Period of Observation:		
From	Dec-02	Feb-05
To	Nov-05	Jan-08

JUMP DIFFUSION

While it is possible to measure stock price volatility under the assumption of continuous time, investors are generally interested in discrete time periods – specifically the expected prices at the times on which they buy and sell the investment. And because the determination of all value unavoidably involves a prospective future-oriented investigation, business valuators are also necessarily interested in the long-term forward estimate of an asset's volatility – generally stated at an annualized rate. The frequency at which an asset price change is measured could conceivably be as small as a minute or as large as a year-over-year change. In attempting to

[9] Some little caution needs to be exercised here in that the comparison of these two betas is not for simultaneous time periods. Instead of using contemporaneous periods, the process has been one of maximizing R-Squared in an attempt to obtain the best fit of market to equity covariance. Such a process is often necessary in order to minimize the distortion caused by non-systematic risk in the equity price samples. In both regressions the standard error of the Y estimate is quite low compared with the least-squares beta.

estimate a long-term, mean reverting average, however, dissecting the changes into periods as finite as a minute is pointless whereas using a year-over-year change with only two data points is far too coarse and therefore subject to statistical error. In the 252 rolling daily graphs presented above we have used the day-over-day closing stock prices in order to calculate the annualized volatility. In some cases we have observed highly significant "jumps" in the one day data. These typically happen because of some company-specific event or news release that the market had not anticipated. These cause a one-time dramatic change (the jump – note that a price change either up or down will always cause volatility to jump upwards) to the firm's stock price and immediately there afterwards the normal degree of volatility is resumed. Note however, because we are looking at a 252[10] day progressive rolling average, a one-time jump has the effect of ratcheting the entire average up for the full 251 subsequent trading days after which time the degree of volatility plummets down to the original level once the triggering event becomes older than 252 days. A visual emphasis of these jumps will improve comprehension:

[10] Normally, there are 252 trading days in any given calendar year.

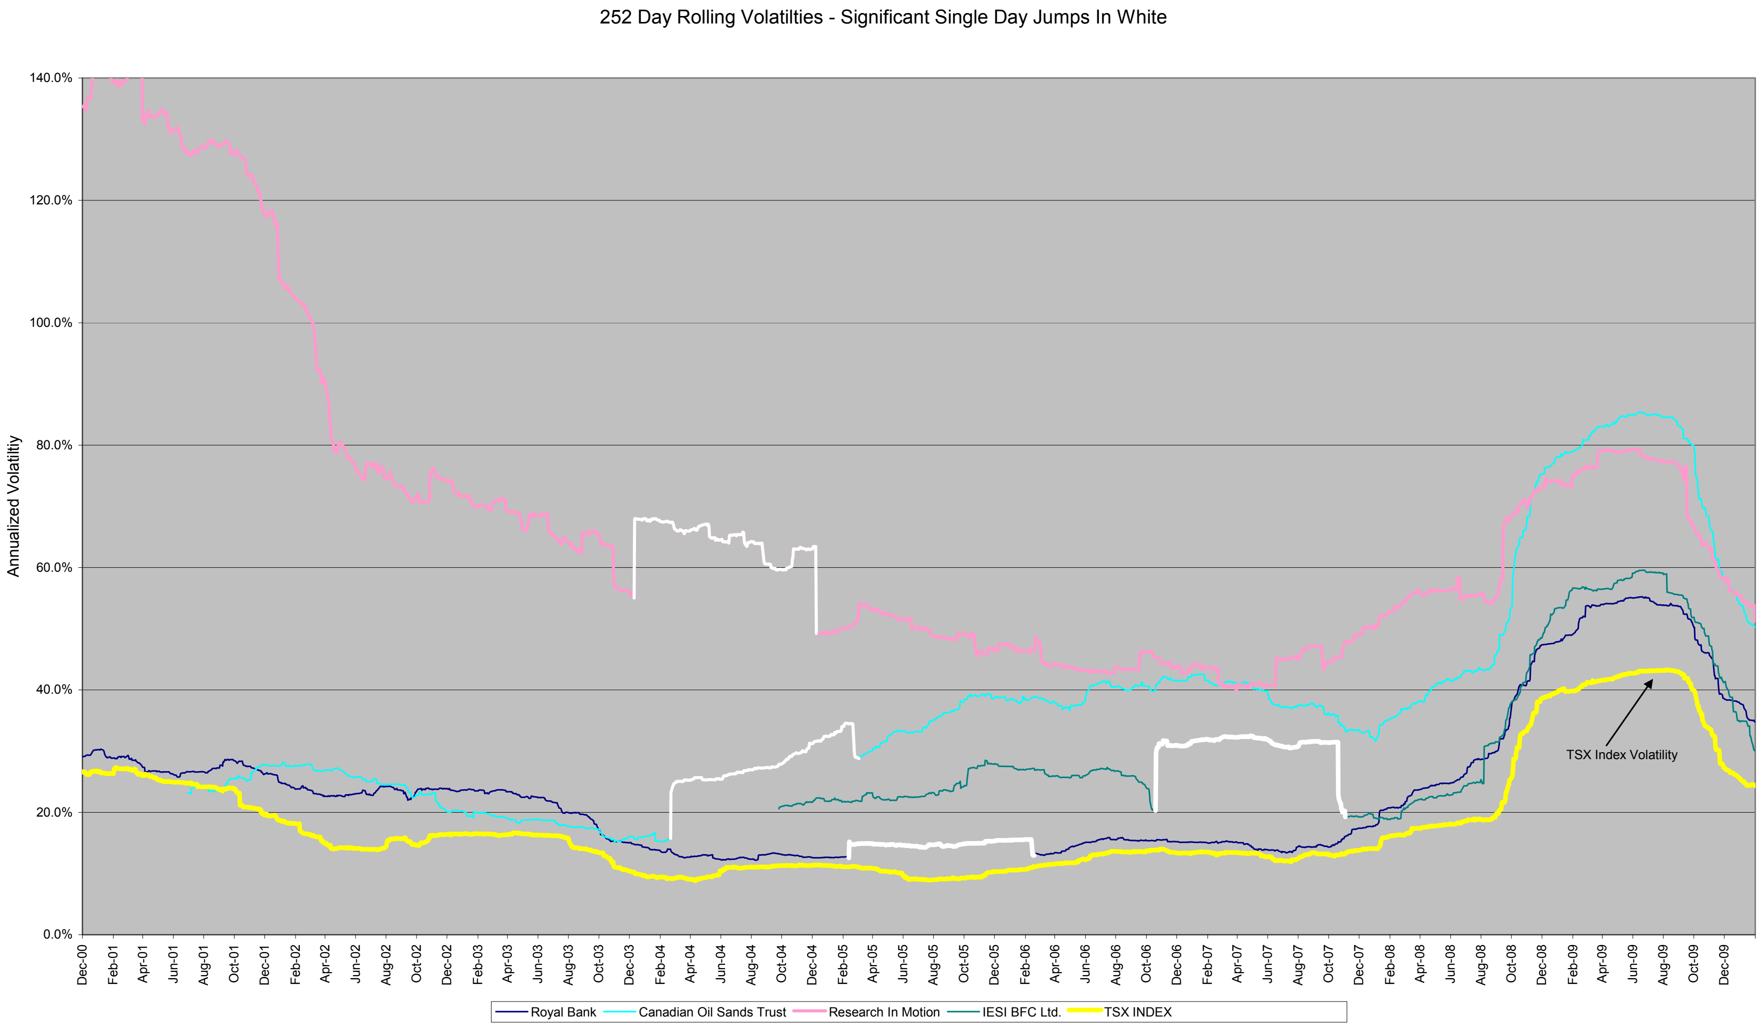
252 Day Rolling Volatilties - Significant Single Day Jumps In White
Annualized Volatility
140.0%
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
Dec-00 Feb-01 Apr-01 Jun-01 Aug-01 Oct-01 Dec-01 Feb-02 Apr-02 Jun-02 Aug-02 Oct-02 Dec-02 Feb-03 Apr-03 Jun-03 Aug-03 Oct-03 Dec-03 Feb-04 Apr-04 Jun-04 Aug-04 Oct-04 Dec-04 Feb-05 Apr-05 Jun-05 Aug-05 Oct-05 Dec-05 Feb-06 Apr-06 Jun-06 Aug-06 Oct-06 Dec-06 Feb-07 Apr-07 Jun-07 Aug-07 Oct-07 Dec-07 Feb-08 Apr-08 Jun-08 Aug-08 Oct-08 Dec-08 Feb-09 Apr-09 Jun-09 Aug-09 Oct-09 Dec-09
TSX Index Volatility
Royal Bank
Canadian Oil Sands Trust
Research In Motion
IESI BFC Ltd.
TSX INDEX

The apparent causes of these abnormal single-day price changes are:

- Royal Bank: Feb. 25, 2005 a 9% single day price gain relating to rejection of a proposal to close subsidiaries in tax havens, amongst other issues presented at the AGM.
- Canadian Oil Sands: March 5, 2004 single day price decrease of 16% relating to the prior day news release disclosing a $2.1 billion capital cost increase in Syncrude's ongoing expansion project.
- Research In Motion: Dec. 23, 2003 single day price appreciation of 50% in relation to prior day news release disclosing that year-over-year revenue for 3rd quarter was up 107%
- IESI BFC: Nov. 1, 2006 single day price decrease of 19% appears to be in response to the Oct. 31, 2006 Canadian Minister of Finance's announcement that income trusts would essentially be taxed as corporations beginning 2011 (at that time IESI BFC was operating as BFI Income Fund … curiously, Canadian Oil Sands, also an income trust, did not show any unusual impact to this announcement, but this may have more to do with the tax-deductible capital intensity of the latter business, and the difference in the amount of tax-free distributions being paid out).

These one-time jumps beg the question "Should these abnormalities be eradicated or otherwise 'smoothed' out of an equity's historic volatility?" The issue is akin to the normalization of historic financial statements. Because the goal is to estimate future performance (in this case, future volatility), any one-time only historic transaction should be eliminated with the belief that it is highly unlikely to ever occur again. The difficulty comes in attempting to assess how confident one is that these events will never reoccur. After all, capital budget constraints do often become violated, revenue targets are often missed (usually on the low side) and taxing authorities have almost unilateral power to restate which entities fall under which provisions.

One approach that does seem reasonable under the circumstances is that, should a valuator decide to eliminate/smooth the impact of a one-time event for the purpose of improving the predictive quality of history volatility data, then the same elimination should also be made for the beta regression data.

STOCK PRICE VOLATILTY & BETA – IS THERE ANY RELATION?

It is tempting to believe that there is some proportionate relationship between the historic volatility of a given stock and its theoretical beta. For example, if the measured average volatility of a specific stock is 20% per annum[11] and this exactly agrees with the volatility of the market index, then it seems logical for that stock to expect to earn the Equity Risk Premium (ERP) in the same proportion that the Market Index would. Ergo, a firm with the same volatility of the market might be expected to have the same beta as the market (i.e. 1.0)[12]. Similarly, it is reasonable to think that, if a particular firm had a long-term average volatility of 30%, whereas the Market Index was 20%, then that firm should earn an ERP of 50% greater than the Market Index (i.e., have a beta of 1.5). This reasoning, however, is faulty.

The reason why there is not a direct proportionate relationship between volatility and stock price beta is multifold. The first, and perhaps most important reason for this, is that the actual observed volatility on an individual stock will include both the systematic (i.e. market-wide) and non-systematic (i.e. asset specific) price variances. While this is also true of the observed stock price beta, theoretically beta should only include the systematic risks[13]. While the volatility of a Market Index probability does contain some small element of the non-systematic price variances of its component individual stocks, it is likely that this effect is greatly mitigated. This is because the more individual stock price movements are consolidated into one composite index; the more likely the individual non-systematic impact of the individual stock prices will cancel

[11] Assuming that the volatility has been measured using the most appropriate frequency (i.e. hourly, daily, weekly or monthly) and that a sufficient period of history has been included in the sample. See Appendix A

[12] After all, if an investor invests in the Market Portfolio, his risk-return expectations would be directly proportionate with the volatility of the Market Index, therefore, by the principal of economic substitution, a single stock that exactly mimics that volatility might be expected to offer generate the same ERP.

[13] See, for example, Accession Capital's paper "Sensitivity of a Stock Beta to Non-Systematic Price Impacts" for an explanation as to why non-systematic risk will distort the measurement of Beta and how sensitive this measurement process can be to even very small amounts of non-systematic price variances.

each other out[14]. The non-systematic distortion of a composite index of 30 stocks can be expected to be exponentially less than that of just one stock. The non-systematic distortion in a composite of 300 stocks can be expected to be virtually zero[15].

When observing the historic price movements of a single stock, the inherent non-systematic price shifts cannot be separated apart from the systematic price movements. Therefore, the empirical volatility of a single stock price history can be expected to be highly distorted by non-systematic price shifts. Conversely, the comparable volatility of the Market Index over the same period of time would be expected to contain very little non-systematic price variances. Ergo, the simple ratio of one security's volatility to that of the Market cannot be expected to simulate an approximation for the beta of that security. Beta, it should be remembered, is the covariance of an individual security's price movements with that of the Market Index. It is often the case when attempting to empirically measure a stock price beta that the process fails because of the distortion caused from non-systematic price movements. Therefore, it is not surprising that non-systematic distortion would also be present in the volatility of an individual stock history and this would also preclude us from finding an alternate means of approximating the true underlying beta of that security. The fact that an individual security's volatility includes non-systematic risk, whereas the Market Index's does not, should not, however, cause us to entirely dismiss the importance of this ratio – as will subsequently be discussed.

Theoretically beta only measures the systematic risk inherent in the target stock price movements. Therefore there are specific firms, and even entire industries that will, on the one hand be relatively insensitive to Market-based risks, but very sensitive to company specific operating risks. These firms can be expected to report relatively low Betas (assuming that the non-systematic distortions obscuring the "true" Beta in the historic data can be consolidated out of the sample) and yet quite high volatility. A good example would be the gold mining industry. Gold Exploration and Production firms are relatively insensitive to market-wide impacts and

[14] See Accession Capital paper "Measuring the Error of Estimation in Grouped Stock Price Betas"
[15] The standard error of the Market Index regression can be expected to decline proportionately with the square root of the number of firms in the index. Therefore, an Index with 300 firms in it would be expected to have only 300^-0.5 = 5.8% the standard error of a one-firm regression. And, a 1000 firm index only 1,000^-0.5 = 3.2%

typically report low betas. Conversely they face very high operational risks (in terms of capital intensity; environmental risks; political/geographic risks; as well as commodity price risks). As a result, in contrast to their low betas, they generally show quite high stock price volatility.

There is another significant difference between beta, as an overall measure of asset risk vs. volatility. By implication Beta provides an approximate measure of 'drift'. That is, if stock price movements are functions of Brownian motion that exhibit tendencies of positive 'drift' or investment growth (and, in the long-run, equity markets as a whole must always be assumed to demonstrate positive drift, or else investors would not be motivated to put their capital at risk), then an individual stock beta also inadvertently provides an indication of drift. If, for example, over the very long term stock markets have been shown to have an average positive drift of 10% per annum, then a stock with a beta of 1.0 should also be expected to experience this same average rate of positive drift.

Volatility, on the other hand, conveys no similar implication of drift[16]. It is possible, for example, for two securities to report precisely the same historic volatility metric, but have entirely differing rates of drift. For example, the following graph shows two stocks (A and B) plotted over a ten period history. Both have a 10% volatility, but Stock A has a continuously compounding average drift rate of 0.9531% per period (or 10% growth over the entire 10 periods) whereas Stock B has an average drift rate of 0.0% per period and 'matures' at the same $100 value that it began with:

[16] In a continuous GARCH model, the rate of drift would be measured by $\mu - \sigma^2/2$ per unit of time (i.e. as $t \rightarrow 0$). In a discrete time GARCH model, the rate of drift μ and variance σ^2 are specific to the time period being measured. In contrast, the Statistics 101 plain Standard Deviation measure has no means of conveying drift.

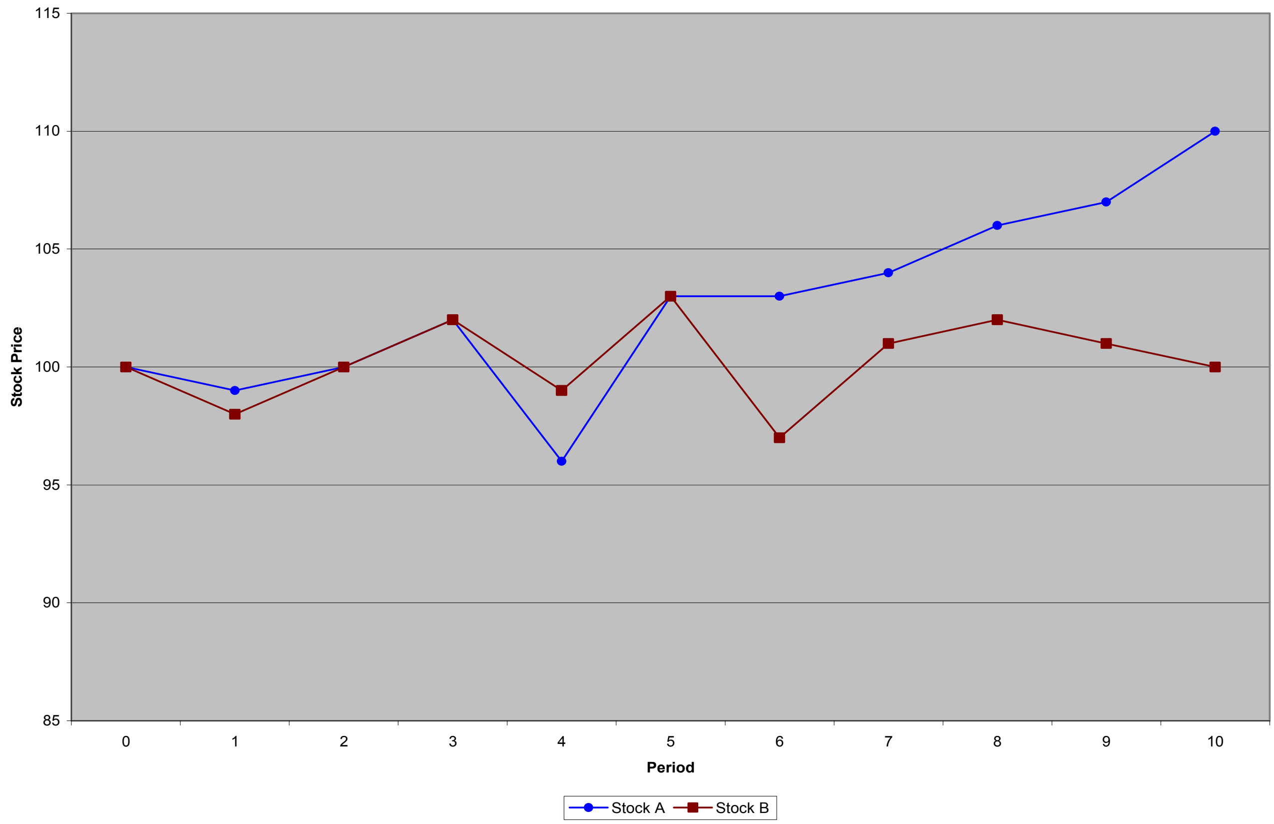

To the typical investor, "risk" is almost always unidirectional[17]. That is, for an investor who is long on the asset, risk is perceived to be incurred as the asset price falls and the opposite is true for the investor who is short the asset. Stock price volatility reports all variances away from the sample mean as an absolute measure – regardless of whether those deviations are positive or negative. Because the plain vanilla measure of stock price volatility fails to provide any indication of drift, it is a less descriptive indicator of asset risk.

Lack of ability to capture the rate and direction of drift is a primary reason why the simple measure of stock volatility cannot be used as an absolute comparator of risk. An equity with a 40% volatility is not necessarily twice as risky as one with only 20%. Volatility is an ordinal rather than cardinal indicator of risk. That is, all other things being equal, an asset with a higher

[17] The exception to this general rule would be a collar or some exotic barrier options where the investor bears no risk as long as volatility is contained within a predetermined range, but a liability is incurred once that range is exceeded.

degree of expected volatility can be said to be a riskier asset – but the quantification of that additional risk will not be discovered by simple standard deviation alone[18].

CONCLUSIONS

An asset's price volatility history can provide a useful insight into the overall risk-assessment of that asset. The information will be more qualitative rather than quantitative. For example, there appears to be no way to deduce the beta of an asset from comparing the asset's volatility to that of the Market Index. Nonetheless, visually inspecting the asset's volatility history against that of the Market Index and other comparable assets should provide the business valuator with important clues of relative riskiness and may induce a more inspired investigation than otherwise would have been possible. In a time, for example, when market volatility is increasing it is very telling that a given security would show signs of decreasing volatility. Similarly, whether or not the target firm has a volatility profile that mimics its peer group is an important question for the business valuator to ask. Reasons for differences may relate to issues of capital structure, the competitive advantages of its product mix, superior/inferior management, tax advantages, differences in operating leverage or intangible assets or future growth potential, etc. A visual representation of the asset's historic volatility is not going to identify which of these differences is the case at hand. However, it is highly likely that if a target firm does differ, in a meaningful way, from its competitors evidence of that difference will visible in the volatility profile.

[18] For a visual representation of the comparative risks between two securities, see Appendix B. When the expected returns and volatility of two securities differ, it is not possible to rank the riskiness of the two simply based upon volatility alone. In these situations, the Sharpe Ratio provides insight into which investment provides the best Risk-Reward mastrix, see the Accession papger entitled: "Approximating the minimum Cost of Equity for Junior Oil and Gas Firms during times of high commodity price uncertainty"

REFERENCES

[1] Andersson, Kristina, *Stochastic Volatility*, (2003), Department of Mathematics, Uppsala University

[2] Bali, T., Demirtas, K., *Testing Mean Reversion in Stock Market Volatility*, (2006), City University of New York

[3] David, A., Veronesi, P., *What Ties Return Volatilities to Price Valuations and Fundamentals?,* (2009), University of Calgary & University of Chicago, SSRN # 1513359

[4] Damodaran, Aswath, *Estimating Risk Parameters,* Stern School of Business, New York University

[5] Engle, Robert, *GARCH 101*, working paper Engle prepared as Visiting Professor of Finance, Stern School of Business, New York

[6] Exley, J., Mehta, S., Smith, A., *Mean Reversion,* Draft of June 2004 presented to Faculty & Institute of Actuaries; Finance and Investment Conference, Brussels

[7] Shephard, Neil, *Stochastic Volatility,* Nuffield College, University of Oxford

APPENDIX A – SELECTING THE FREQUENCY AND TERM OF THE VOLATILITY SAMPLE

In the preceding paper the observations of stock price volatility were based upon day-over-day total changes in return (i.e. closing stock price differences, adjusted for dividends paid) and the term of each sample was 252 trading days. These results were then annualized[19] by multiplying the 251 daily observations by √252.

The questions remain, 'Why use daily, as opposed to hourly, weekly or monthly measurement frequencies?' and 'Why use a sample term of 252 trading days?'

Unfortunately, these are questions that mathematical science will not provide us with any hard and fast rules for. The frequency (i.e. the length of time occurring in between each discreet returns observation) should be dictated by the purpose of the estimate. For example, an arbitrager who is attempting to detect the potential mispricing of a security for a day-trade will most certainly wish to use a frequency of an hour (or less if the market is active enough). However, for the business valuator, attempting to determine the cost of capital on an asset that has several decades of economic lifespan ahead of it, there is absolutely no point in using a frequency of one hour.

The shorter the frequency, the greater the likelihood that more non-systematic price shifts will be captured in the sample data. In avoidance of this problem one might decrease the frequency, perhaps to month-over-month intervals, but this too introduces distortions. As the standard deviation denominator (n – 1) gets smaller, the statistical reliably of the results decreases. In turn, in order to avoid this inaccuracy, one might increase the Term of the sample (i.e. use 120 months of historical observations, rather than 252 days from the past trading year). Unfortunately, there is a trade-off with this approach as well. Increasing the Term of the sample may begin to incorporate volatility data that is no longer representative of the future expectations for the firm. For example, in the case of RIM, including the Dec 2000 to Dec 2003 volatility data in the sample would significantly increase the overall long-term average – but it is unlikely that RIM's past is representative of its foreseeable future (barring the effects of a world-wide financial crisis that skewed all volatilities between Sept 2008 and Feb 2010).

As an indication of how increasing the Term of the sample can dramatically smooth the deviations away from the long-term mean, the following graph shows a term of 1,260 days (i.e. 5 trading years), still using a daily closing price frequency. Note that now, with a denominator of (n – 1) = 1,259, the impacts of the single day jumps (see Jump Diffusion section, above), are virtually undetectable.

---

[19] "The standard deviation of the proportional change in the stock price in a small interval of time Δt is $\sigma\sqrt{\Delta t}$. John C. Hull, *Options Futures, and Other Derivatives, Third Edition* Prentice Hall, 1997, pg. 219

1260 Day Rolling Volatilities - TSX Sector Leaders

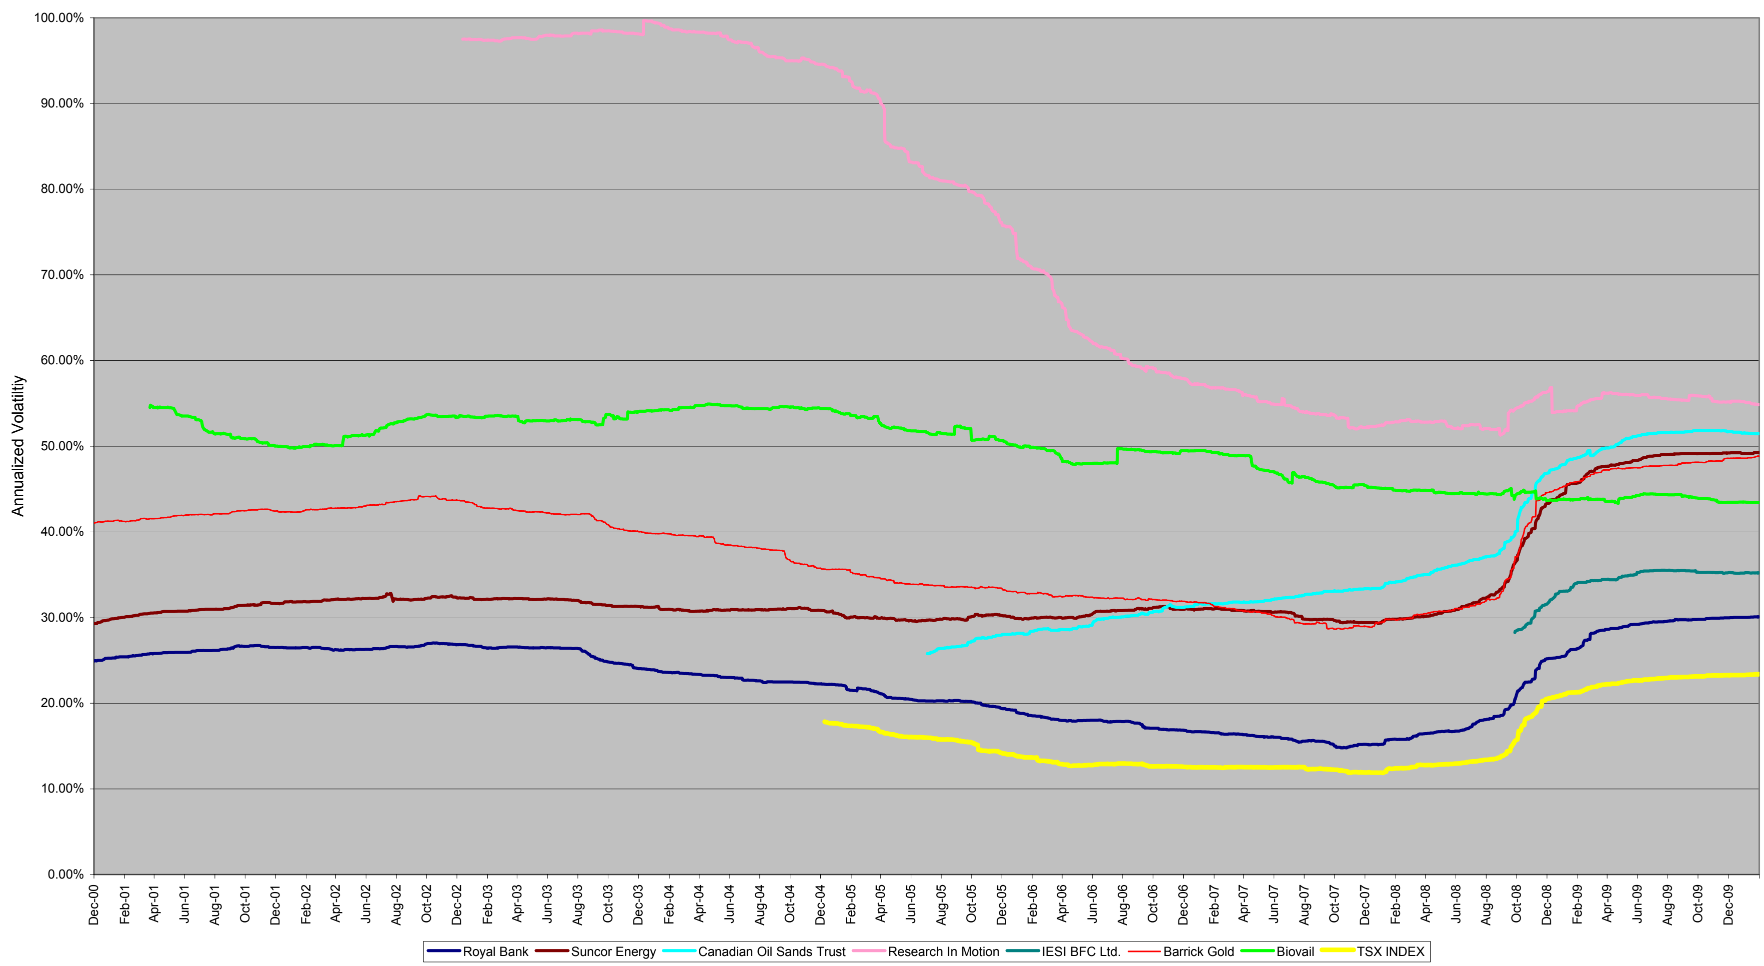

APPENDIX B – GRAPHIC REPRESENTATION OF RETURNS & RELATIVE VOLATILTIY

The standard deviation/volatility of a firm's total daily returns (i.e. the day-over-day change in dividend-adjusted price) fails to indicate the historic drift or rate of growth/decay in the stock price. This shortcoming can be compensated for, somewhat, by visually plotting the frequency histogram of those daily returns.

It will be prudent to use some theoretical examples to cover the basic concepts. For example, if stocks were continuously traded, it would therefore be possible to obtain a continuous frequency distribution of those returns. And, if two similar stocks had normally distributed returns (meaning, a typical bell-shaped probability distribution, where SKEW[20] = 0 and therefore the two halves of the curve can be expected to be perfectly symmetrical), then a visual comparison of the two distributions would be revealing:

SCENARIO ONE

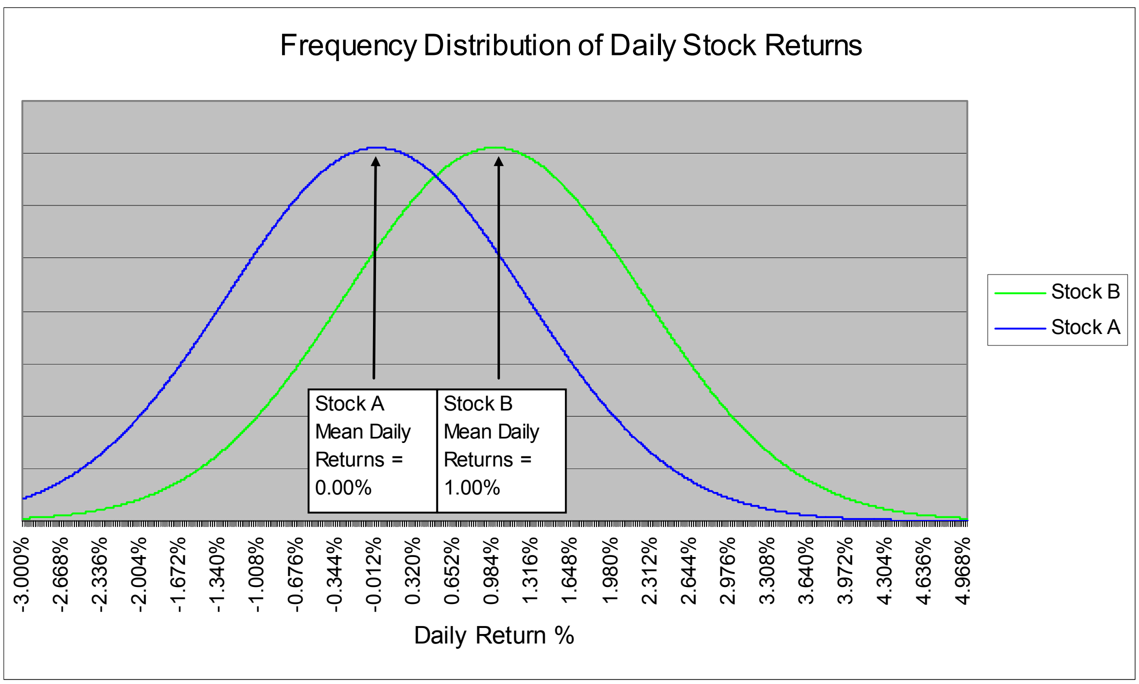

In the graph above the daily return percentages for Stocks A and B have been plotted on the X axis, with the frequency, or number of occurrences of each daily return represented on the Y axis. The most frequently occurring return is represented by the peak of each distribution and is

[20] The definition and importance of SKEW will be discussed later.

known as the mean or average or Expected Return[21] (the first moment of the integral according to the fundamental theorem of calculus). So, visually comparing the two stocks A and B, it can quickly be ascertained that the relative risk of the two firms are identical – that is, the shape and dispersion of the two curves are both the same. This means that the volatility, or standard deviation of the historic daily returns is also identical (mathematically constructed, in this case, that both curves has an annualized standard deviation of 20%).

Moreover, we can clearly see that Stock B dominates Stock A. On average Stock A has generated an average daily return of 0.0% whereas Stock B is 1.0% (In order to make the separation between the two curves graphically more distinct, Stock B was given a 1.0% daily return. In reality, such performance could not be expected to last long, as this equates to an annual return in excess of 1100%). The scenario is obviously hypothetical. In the real world, all the A Investors would realize that Stock B offers a much superior return with exactly the same risk exposure and would be inclined to sell-off their A holdings until the price of A had been bid down to the point where the long-term expected returns exactly equaled that of B.

A much more realistic representation of actual daily return distributions would be:

SCENARIO TWO

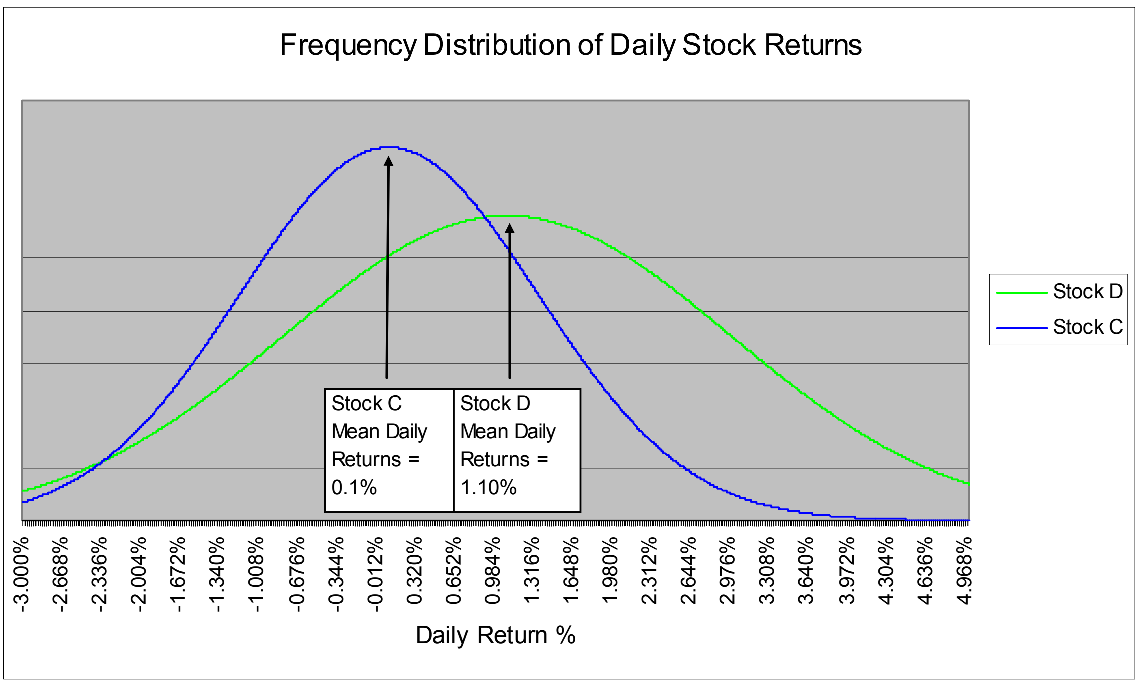

[21] The presumption in analyzing historic data, such as daily returns, is that they provide some indication of what is expected in the future. When this is clearly NOT the case (e.g. RIM pre-2000 stock performance compared with post 2010 expectations), there may be no benefit to making any inferences from historic data. If the future performance is expected to vary for known and reasonably predictable events, it still may be useful to begin the analysis from historic data, and alter the future outcomes accordingly.

Now we can see that, while Stock C still has normal distribution with an annual volatility of 20%, but the average historic daily return was 0.1% (as no investor would voluntarily purchase an equity where the long-term expected return was 0.0%). And, similar to the first scenario, the differential between the Expected Returns is still 1.0%, as the mean of Stock D is now 1.1%. Now, however, the two equities can no longer be said to show the same risk exposure. Stock D has an annual volatility rate of 30% and this is evidenced by the much wider dispersion of the D curve. Also, the frequency (height on the Y axis) at which Stock D has earned a 1.1% return in the past is less than how often Stock C has earned a 0.1% return. Now it is no longer apparent if D dominates C or the other way around[22]. As is typical, the stock showing the higher expected return, Stock D, also demonstrates a higher degree of risk as represented by the expected volatility.

SKEW OF A PROBABILITY DISTRIBUTION

Skew is the measure of asymmetry in a distribution. By definition then, a skewed distribution is not normal (i.e. does not conform to the standard Gaussian bell-shaped curve). A left-skewed distribution has either more or quantitatively larger negative deviations [e.g. $(y_i - \hat{y})$] than positive. And, a right-skewed distribution will have more/greater positive deviations. Suggestively, then, when measuring the historic outcome of daily returns we would prefer to have a positively skewed distribution. This would be evidence that the frequency or size of returns above the mean return is higher than below it. Note, however, that a positive skew does not imply a positive mean. It is still possible to have a positively skewed distribution where the mean was either zero or negative.

Practically speaking, knowing the skew of the returns deviations is of little use. This is because the directional impact of the skew is already reflected in the historic mean return. It is, for example, far more important to know that a security has generated (and possibly would be expected to continue to generate) a 15.0% annual return than being told its skew statistic is -0.40. Moreover, an investment with a 15.0% annual return and -0.40 skew should generally be considered a far superior investment than one with a 2.0% return and 3.4 skew. Any asymmetry in a historic distribution is already reflected in the mean return and the skew statistic does not provide any further predictive quality. That is, the 2.0% mean return is already the result of the highly positively skewed outcomes and unless there is some reason to believe that the skew will become even more positively pronounced, it would still be much more desirable to hold the 15.0% return investment with the negative skew (depending, of course, upon the relative riskiness/volatilities of the two alternatives).

Two commonly used formulas for Skewness, in a discreet environment is:

$$\sum(y_i - \hat{y})^3 / ((n-1)\sigma^3)$$

[22] It would be possible to approximately quantify the change in overall risk from Scenario One to Scenario Two by calculating the differences in the areas between the intersection of A∩B and C∩D, but perhaps this is a topic left to another paper.

Or

$$[\sum(y_i - \hat{y}) / \sigma]^3 \cdot (n / ((n-1)(n-2))$$

Where,

y_i is the single day total return instance

$\hat{y}$ is the sample average of all the n number of daily returns

n is the number of daily returns in the sample

σ is the standard deviation of the daily returns

and, of course, $(y_i - \hat{y})$ is deviation from the mean for the "i"th observation. Therefore the outcome of the Skew statistic depends upon whether the either majority of the $(y_i - \hat{y})$ deviations are positive, or, whether the magnitude of the positive $(y_i - \hat{y})$ deviations outweigh those that are negative.

Putting this into context, if we examine the relative skewness of the seven major firms referenced in the body of the paper:

JAN 2000 – AUG 2008 SAMPLE

SECTOR	NAME	SYMBOL	Average Volatility 252 day rolling	# of Observations	Mean of Daily Returns - Annualized	Standard Dev. of sample of observations (annualized)	Mean Return per Standard Deviation	SKEW Statistic
Mining	Barrick Gold Corp	ABX	34.5%	2,177	4.71%	34.97%	0.13	0.081
Clean Tech	IESI BFC Ltd.	BIN	25.7%	981	2.93%	28.22%	0.10	-2.993
Life Science	Biovail	BVF	49.0%	2,177	-9.29%	49.46%	(0.19)	-1.203
Income Trst	Canadian Oil Sands Trust	COS.UN	30.7%	1,795	36.80%	33.33%	1.10	-0.437
Technology	Research In Motion	RIM	73.5%	2,177	32.85%	78.14%	0.42	-0.298
Integrated	Royal Bank of Canada	RY	19.8%	2,177	16.09%	21.39%	0.75	0.265
Energy	Suncor	SU	31.2%	2,177	26.37%	32.32%	0.82	-0.118

We note that the three firms with the largest historic returns over the period (COS.UN, RIM and SU) all show a negative skew. In fact, COS.UN reports the highest risk-adjusted return of the entire group (being 1.10% of annual return for each unit of volatility) and still experienced a negative skew of -0.437

The skew statistic can be thought of as the ratio of standard deviations that lie above the distribution mean compare with those that lie below it. In a "normal" Gaussian distribution, the ratio is one-to-one and hence the left side of the curve is perfectly symmetrical with the right. However, even in the case where the skew statistic is showing significant asymmetry, it provides little insight towards assessing the overall riskiness of the asset compared with the overriding importance of the distribution mean.