Conn Valuation Services Ltd.

RISK QUANTIFICATION: THE SEMI VARIANCE AND OTHER RISK MEASURES

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

Remember those optical illusions where they ask you to visually identify which line, A or B is longer?

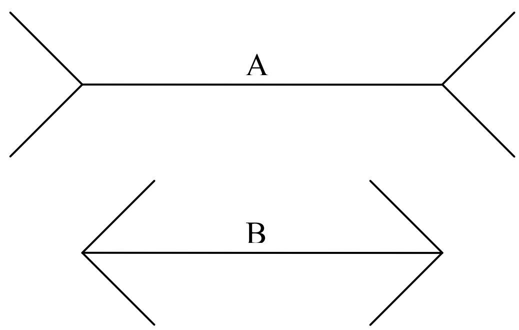

… and then surprise you with the disclosure that both lines are exactly the same length. Well, it turns out that financial risk quantification is much the same: superficial observation of risk will often deceive us.

For example: Our client tells us he has been approached by two different entrepreneurs seeking to sell out their position in two different projects (each, conveniently, with a 10 year lifespan). The client wants your opinion as to which of the expected net free cash flows is more risky (Project A or Project B):

TABLE A

$000's	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10
Proj A	50	350	(350)	(200)	300	400	1,300	450	300	(450)
Proj B	200	(150)	500	(200)	200	225	275	225	175	125

THE ANNUAL CASH FLOW ESTIMATE

Perhaps, before discussing the comparison between the two independent projects, we should consider the conceptual issues behind the construction of any single net cash flow. We will use Project A as an example. Expected cash flow projections are represented as a sequence of annual point estimates. But in reality each of these annual amounts reflects the most likely or "expected" cash inflow or outflow from a contingent range of possible outcomes. We can, therefore, visualize the Project A cash flow as a sequence of ten normal distributions[1] (i.e. the Gaussian bell curve) wherein the mean or average of the distribution becomes the point estimate for that year. In the following graph, I have turned the annual standard curves on their sides so that time, in years, can be represented by the X axis:

GRAPH 1.

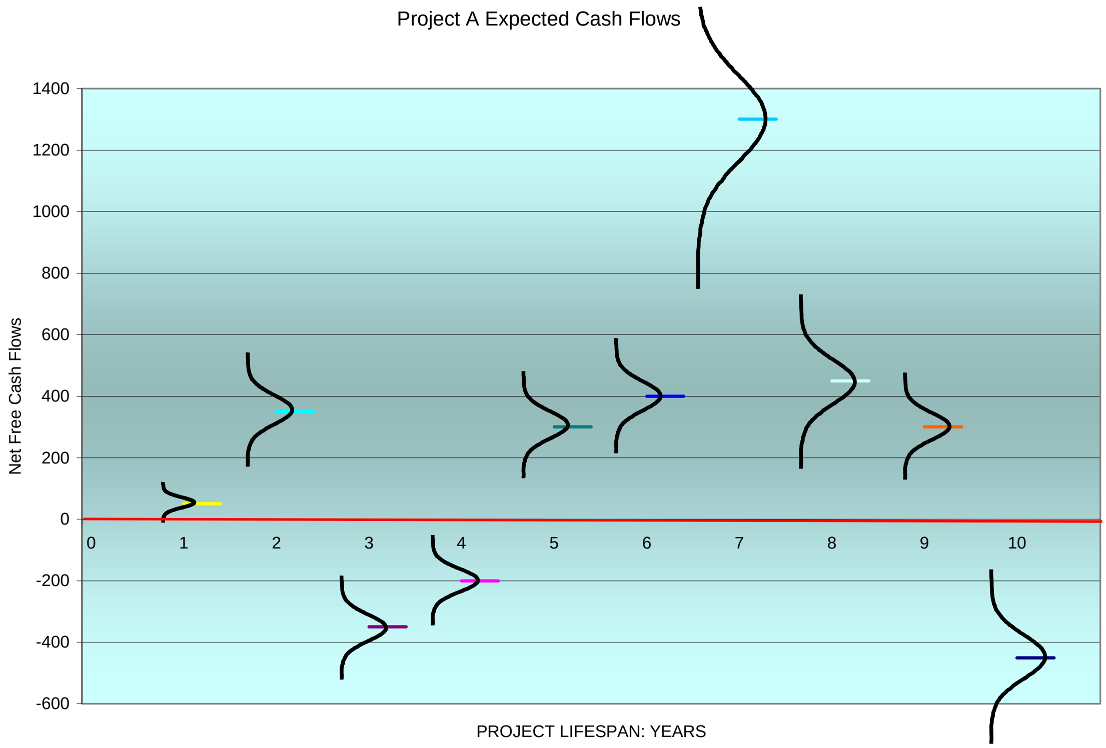

What should become apparent just by looking at the graph is that each year individually is an exercise in probability. And, unless we are dealing with risk-free securities, there is never any assurance that the actual cash outcome will equate to the expected mean. The estimator's confidence will, however, be reflected in the dispersion of the distribution. So, at the end of Year One[2], Project A is expected to generate \$50K of positive net free cash. If either one Estimator (who has the requisite expertise in this project area) is asked to predict the cash outcome of this year 100 times, or, alternatively, one-hundred equally qualified Estimators are asked to predict Year One once, what we would find is that sample of 100 estimations would have a measurable standard

[1] It is conceivable that the annual estimates are represented by non-normal distributions. For example, Estimators who are pessimistically biased against the project could consistently estimate distributions that are skewed to the left whereas optimists would produce right-skewed distributions.
[2] Year-End cash receipts are assumed for convenience.

deviation. The smaller the standard deviation in nominal dollar terms is relative to the mean, the higher the confidence one could have in the accuracy of that prediction[3].

Some may argue that, in actual practice, our client would never go out and engage 100 experts in this business area in order to solicit their opinions about the cash projection. True. But what does customarily happen is that a single expert (or a team of experts) will construct a model that aggregates all the likely Year One Revenues; Operating Expenses; Taxes; CAPEX, etc. in order to produce an estimate of that year's net free cash flow. And, typically, there is a range of perfectly viable inputs for each revenue and expense item that are tested to determine their impact on the final results. It is the output of these sensitivity tests that could easily amount to 100 trials for that year's net free cash flow.

If, for example, we found that the standard deviation from these trials for the Year One estimate was $5K (i.e. 10% of the mean), then we would normally expect that plus-or-minus 3 standard deviations (or, in other words, +/-$15K, by the Empirical Rule) would represent almost 100% of all possible expected outcomes. Therefore, we would have a great deal of confidence in believing that the Year One cash receipt would fall somewhere in the $35K to $65K range. Conversely, if the Year One standard deviation was measured to be $20K (40% of the mean), then the plausible range increases; from $(10)K to $110K. This would inspire much less confidence in the estimate. The relative dispersion of the annual estimates will have an impact upon the appropriate discount rate chosen in a DCF (discounted cash flow) analysis.

It is reasonable to expect that there will be an internal consistency in confidence from year-to-year throughout the cash projection. That is, if the Estimator's Year One confidence reflects a standard deviation that is 10% of the mean (which is approximately the proportions depicted in GRAPH 1. above), then we would presume the same degree of confidence reflected in all the subsequent years. So the Year One range of virtually all possible outcomes is represented by $50K +/- 3($50 x 10%) = $35K ~ $65K, and similarly, in Year Seven: $1,300K +/- 3($1,300 x 10%) = $910K ~ $1,690K. If, for some reason, the Estimator(s) were not able to adhere to a proportionally consistent standard deviation in each year[4], then we would expect that adequate disclosure be made and factor in this change in confidence into the overall risk assessment.

COMPARING TWO DIFFERENT CASH FLOWS

Now that we have considered the annual risk components within each cash flow projection, we can turn our attention to the original question which asked about the relative riskiness between A and B. At first blush, we can see that Project A has a much wider range of projected cash flows than Project B ($1.75M for A [-$450 to $1,300] compared with only $0.7M for B). Range gives some indication of overall risk – except that it will be distorted by outliers. So, next we observe the annual average cash in-flows for A is $215K whereas it is only $158K for B. Surely it must be less

[3] To be convinced of this perspective, just imagine asking 100 financial experts what their prediction of the fixed-rate interest income would be for a given year on a specific treasury bond. In this case the deviation would be zero and (hopefully) all 100 experts would exactly agree on a point estimate for the annual income.
[4] Perhaps the $1,300K estimate is dependent upon the town hosting some festival or special event in Year Seven – but there is less certainty of this occurrence than the predictions of the surrounding years.

risky to have a higher expected average cash flow than a lower one? Well, perhaps. But then we must acknowledge that nominal, non-discounted cash flows are a relatively meaningless financial indicator. Nominal point estimates of dollars projected over a long period of time convey little useful information as to the 'riskiness' of the expected cash flows. Further, nominal dollars do not even capture the time-value-of-money component of the expected receipts. Most of A's large positive cash inflows occur at the later stages of the project, when the risk-and-time-adjusted value of these cash receipts will be much less. We should, of course, discount these annual cash flows so they are represented at their present value and then the comparison could then be made 'apples-to-apples'. But this approach begs the question 'at what risk-adjusted discount rate'?

We can easily calculate the standard deviation of both cash estimates. After all, the equity markets are always representing risk by volatility, and volatility is just the standard deviation of the period-over-period stock price movements. As it turns out, the standard deviation of A's cash flow is $501K whereas B's is $202K [5]. So, Project A must be more risky than B – right? Intuitively the concept of standard deviation as a measure of risk makes sense. In laymen's terms the standard deviation can be thought of as the average size of the expected movements away from the sample mean. So, in Project A's case, we know that, on average, the expected annual return is a positive $215K per year, and the one standard deviation of $501K away from this would indicate a range of (negative) $(286)K to $716K. A movement of within plus or minus one standard deviation occurs about 66% of the time (from the Empirical Rule again - assuming the expected cash flows are normally distributed). It doesn't take a rocket scientist to conclude that such a realm of possibility $215k +/- $501k 66% of the time *must be* more risky that Project B's $158 +/- $202. Right? Well, yes, but there are some caveats.

Greater year-over-year volatility does indicate increased risk. However, the 10 annual project estimates are, in turn, a function of the degree of intra year dispersion. Recall that GRAPH 1 showed Project A annual distributions to have a standard deviation that was approximately 10% of the mean point estimate. If the Project B annual distributions has this same degree of confidence/dispersion, then, ceteris paribus, the wider volatility in Project A's year-over-year cash estimates would indicate a higher level of risk. On the other hand, if the degree of dispersion within Project B's annual cash estimates is much greater – for example, if the standard deviation of a each year's point estimate was 40% of the mean – then a much lower degree of confidence would be had about Project B's point estimates compared with Project A. In such a case A and B's cash flow estimates are not directly comparable and it does not necessarily follow that the greater year-over-year volatility of A is proof of greater risk.

Note that the two projects differ intrinsically in that B, once purchased, is expected to maintain a cumulative net nominal positive cash balance. In contrast, Project A will require additional capital injection in Year 4 if the subsequent positive cash inflows are expected to materialize. Further, A requires an additional net outflow in Year Ten (perhaps this is a remediation charge or contractual commitment to restore assets to their previous condition).

As is often the case in the process of risk quantification, there is an element of circular reasoning that comes into play. For example, if we knew, independently – from sources external to these specific cash projections, what the FMV was of each project (a market approach), then we could

[5] Both the $501K and $202K standard deviations assume nine degrees of freedom.

determine the discount rates that would equate the expected cash flows with that FMV purchase price. Conversely, if we knew from an exogenous source, what the appropriate risk-adjusted discount rate was for each expected cash flow (an income approach), then, of course, it would be easy to determine the FMV for each project. As it stands, we have neither of these external indicators. However, we can apply some inductive statistical reasoning.

THE SEMI-VARIANCE

One of the shortcomings of standard deviation as an indicator of risk is that it makes no distinction between volatility direction. That is, positive deviations above the mean carry the same weight in the standard deviation formula as the negative below. In 1952 Arthur Roy and Harry Markowitz independently began writing about the Semi-Variance and by 1959 Markowitz had dedicated an entire chapter to the subject in his seminal book "Portfolio Selection"[6]. Since that time, a great deal of academic research and empirical investigation has been dedicated to the idea of 'downside' risk quantification.

In its simplest form, the idea is that investors will be more concerned about the potential for volatility that leads to losses and erosion of capital than the variability that might occur when returns are positive. The Semi-Variance (SV) formula is:

$$SV = E[\min(x_i - b, 0)]^2 \qquad (1)$$

Where:

$E \equiv$ The Expected Value (indicating the Mean or Weighted Average)

$\min \equiv$ Minimum of two outcomes

$x_i \equiv$ the "i" th observed value of x

$b \equiv$ could be any desired investment threshold or 'target return'[7], however, we are going to confine ourselves to "b" as the Cash Flow Mean

Therefore:

$$SV = (1/(n-1))\sum(\min(x_i - \text{Mean}, 0))^2 \qquad (2)$$

Where:

$n \equiv$ number of observations (and we have assumed that this is a sample, so there is one less degree of freedom than observations)

In Project A's case, for example, the $\sum(\min(x_i - \text{Mean}, 0))^2$ part of the formula is represented by:

[6] Markowitz, Harry M. "Portfolio Selection" (1st Edition). New York, John Wiley and Sons (1959)
[7] See, for example, Jun Kim & David Wallace (1998) "Mean-Semivariance Analysis: Risk and Opportunity"

$$(\min(50 - 215,0))^2 + (\min(350 - 215,0))^2 + \ldots + (\min(-450 - 215,0))^2$$

Note as well that the Semi-Variance can easily be rearranged to provide the extent of variance above a given threshold:

$$SV = (1/(n - 1))\sum(\max(x_i - \text{Mean},0))^2 \qquad (3)$$

… which, for lack of a better identifier, we could refer to as the measure of 'Upside Risk'.

The Semi-Variances for both projects are:

TABLE A

Project A		Project B	
Downside Semi-Vari	Upside Semi-Vari	Downside Semi-Vari	Upside Semi-Vari
$106,767	$ 144,372	$ 24,824	$ 16,016
42.5%	57.5%	60.8%	39.2%

The dollar values themselves are not a meaningful measure, as they are presented in thousands-of-dollars-squared. But it is the proportion of those observations that gives us greater insight into the nature of each project's level of risk.

In Project A's case, 57% of the variability occurs above the expected mean[8], indicating that the primary threat to these cash expectations would be missing the highs expected in years five through nine. In Project B's case, there is a greater degree of variability below the mean, suggesting a higher degree of uncertainty about years two, four and ten.

THE REQUIRED RATE OF RETURN

Ultimately financial risk will be encapsulated in a required rate of return, or risk-adjusted cost of capital that will reflect the degree of certainty in the annual cash predictions. It would be an aberration of economic logic that the riskier project would ever have the *lower* cost of capital. If, for example, the Market is pricing Project A with a 12% required return and B with only an 8% - then A *must be* the riskier project[9]. All other things equal, a project where the standard deviation is only 10% of each annual expected cash inflow/outflow will have a lower cost of capital than one that has a 40% standard deviation.

[8] Therefore Project A is positively skewed. The larger proportion of variance below the mean in Project B's case indicates negative skew.

[9] This is not to say the Market can never be mistaken or ill-informed or deceived about the nature of various risk components. Based upon the disclosed information, however, the Market will always price the project that it perceives as being riskier with the higher cost of capital.

In the case at hand, we don't know what the Market would say about the required returns for Projects A and B, but it is easy enough for us to speculate across a whole range of likely returns. Note that:

$$PP_a = \sum(x_i \,.\, e^{(-ra \cdot t)}) \qquad (4)$$

Where

$PP_a \equiv$ Purchase Price of Project A

$x_i \equiv$ the "i" th expected cash inflow of Project A (so, x_1 = +$50, x_2 = $+350, x_3 = -$350 …)

e ≡ is the natural log base

ra ≡ is the Project A discount rate applied (stated as a continuously compounded rate, which is computationally easier to work with than an effective annual rate)

t ≡ is the project year (1 through 10) in which x_i is expected to occur

Note that when ra = the risk-adjusted cost of capital, as determined by the Market, then

$$PP_a = FMV_a \qquad (5)$$

The Purchase Price calculated in equation (4) will equal the Fair Market Value of Project A. The distinction between PP_a and FMV_a will be an extremely important one for our client. Similarly, for Project B:

$$PP_b = \sum(y_i \,.\, e^{(-rb \cdot t)}) \qquad (6)$$

Where

$y_i \equiv$ the "i" th expected cash inflow of Project B

Equations (4) and (6) can be rearranged to give

$$\sum(x_i \,.\, e^{(-ra \cdot t)}) - PP_a = 0 = \sum(y_i \,.\, e^{(-rb \cdot t)}) - PP_b \qquad (7)$$

Which, in turn, can be rearranged to:

$$PP_a - PP_b = \sum(x_i \,.\, e^{(-ra \cdot t)}) - \sum(y_i \,.\, e^{(-rb \cdot t)}) \qquad (8)$$

Or, in other words, the difference between the purchase price of A vs. B will simply be the difference in the aggregate discounted value of the ten Project A and B cash flows (using discount rates ra and rb, respectively). Moreover, if ra and rb happen to equate to the Market determined risk adjusted cost of capital, then $PP_a - PP_b = FMV_a - FMV_b$.

For example, if both ra and rb are both speculated to equal 10%(cc), then $PP_a - PP_b$ = \$1,144K - \$894K = \$250K. One would expect to pay \$250K more for Project A than for B. Of course, such speculation is relatively meaningless unless we have reason to believe that 10% is the correct discount rate to apply to both projects. We can, however, gain some insight into the issue by creating a matrix of $PP_a - PP_b$ values over what we might consider the most likely range of ra and rb rates. If we speculate that ra and rb would normally fall somewhere between the 8%(cc) to 20%(cc) range, and then plot equation (8) on a surface map, we get:

GRAPH 2

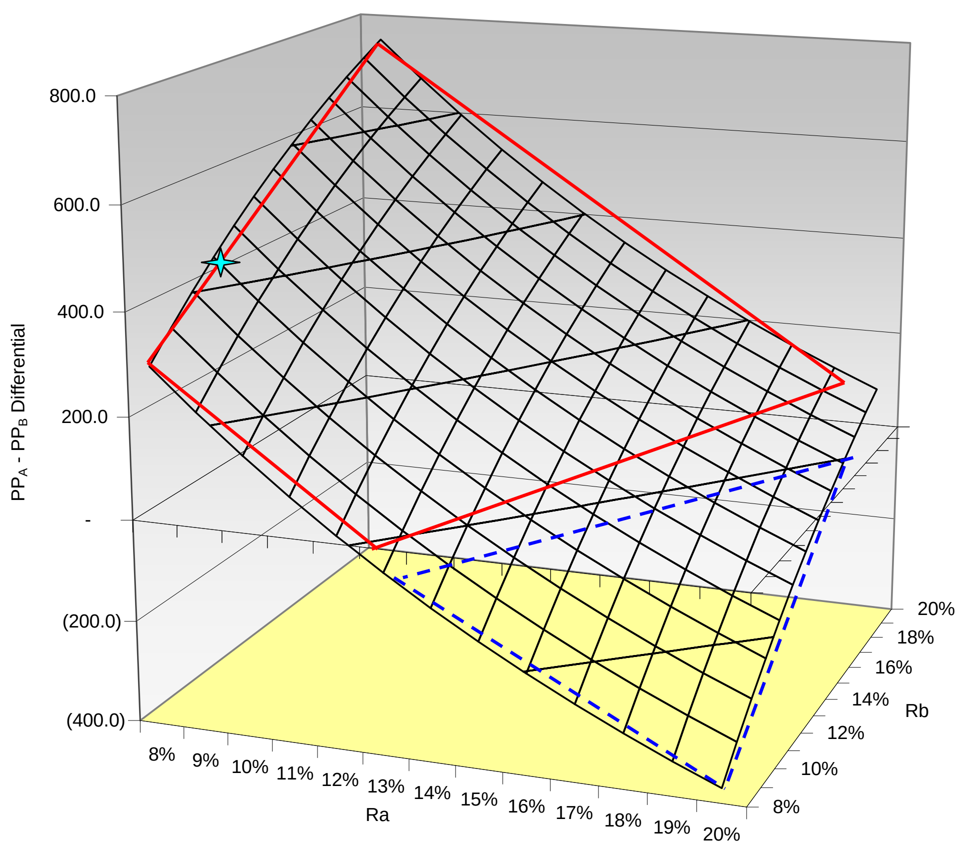

Visually GRAPH 2 tells us that there are a considerably more number of points were Project A will be worth more (i.e. have a higher Purchase Price) than Project B. If, for example, ra = 8.0% and rb

= 11%, we can see on the graph that Project A would be expected to have a purchase price approximately $400K higher. And, this can be confirmed by subtracting equation (6) from (4) to get $1,296 - $848K = $448K. All the combinations of ra vs. rb discount rates wherein Project A is expected to be worth more than Project B is shown above the horizontal zero line and roughly represented by the RED trapezoid in the graph. Similarly, all the combination of discount rates wherein Project B will be worth more than A (i.e. $PP_a - PP_b$ leads to a negative number, meaning B is worth more) is represented by the BLUE dashed-line triangle. And, as is visually apparent, the area of the red trapezoid is much larger than the blue triangle.

Well – so what? The matrix of discounted values and the pretty picture haven't led us any closer to determining what the correct risk-adjusted cost of capital should be. And, on the proviso that the client eventually does pay FMV for either project, then he will be receiving the correct risk-adjusted return for that investment. It doesn't really matter that there are more opportunities for Project A to have a greater absolute FMV than Project B – only that the expected cash returns are risk-appropriate relative to the price paid[10].

These observations are true, but there is one important reason why the client should be very interested in the nature of GRAPH 2. There is probably no single more effective risk mitigator than acquiring a project or asset at less than FMV. If the client is a particularly skilled negotiator, the larger area and steeper curve of the red trapezoid indicates that there is much more upside potential with Project A than Project B (something that the Semi-Variance already indicated). If only one project can be taken on, there is more opportunity to mitigate risk (either by negotiating a below-FMV purchase price or incorporating post-purchase performance concessions with the current vendor) for Project A and still generate a larger absolute dollar return.

THE COST OF BEING WRONG

Financial risk is synonymous with the term "the cost of being wrong". It is equal to the probability of the future cash inflows and outflows being different than what is expected[11]. Since there is a direct correlation between the dispersion in the annual cash flows with the appropriate discount rate, then being wrong about the cash flows means being wrong about the discount rate. Indeed the greatest risk with respect to acquiring Projects A and B would be in getting the discount rate wrong. The downside risk would be in underestimating the discount rate (e.g. paying a purchase price based upon a 10% discount rate only to subsequently learn that 12% was the correct risk rate – making the mistake the other way around would be a pleasant surprise).

We can easily quantify the cost of being wrong into actual dollars for each project and compare the two. For example, using equation (4), we can calculate what the cost of purchasing Project A for the assumed-correct 10% cost of capital, only to shortly learn that 12% would have been the correct rate:

---

[10] Assuming that the client's absolute personal risk tolerance does not max-out somewhere below 20%.

[11] Note that this is a profoundly different concept than measuring the difference between the expected and subsequent ex post actual cash receipts eventually are realized at. Valuation is prospective. Therefore, once actual cash flows are known, the only discount rate that would ever apply would be the risk-free rate.

$$Loss_a = \sum(x_i \cdot e^{(-0.12 \cdot t)}) - \sum(x_i \cdot e^{(-0.10 \cdot t)})$$

$$= \$1{,}011K - \$1{,}144K = -\$133K$$

Such an error would immediately lead to a loss of $133K. That is, the client would have overpaid $133K for Project A. As a rough approximation, since we speculated on a 10% cost of capital but subsequent market evidence proved 12% to be more appropriate, we could say that the cost of being wrong in this case amounted to $133K/2 = $66K per 100 basis points. Continuing this same reasoning, we can determine just how sensitive Projects A and B are to a change in 100 basis points (via formulas (4) and (6)) throughout the reasonable range of 8%(cc) to 20%(cc):

	Project A		Project B	
Risk% cc	PP_a (000's)	Diff per 100bp	PP_b (000's)	Diff per 100bp
8.0%	1,296.1		995.2	
9.0%	1,217.6	78.5	942.9	52.3
10.0%	1,144.1	73.5	894.1	48.9
11.0%	1,075.3	68.8	848.4	45.7
12.0%	1,010.9	64.4	805.7	42.7
13.0%	950.7	60.3	765.7	40.0
14.0%	894.3	56.4	728.3	37.4
15.0%	841.5	52.8	693.3	35.0
16.0%	792.2	49.4	660.4	32.8
17.0%	746.0	46.2	629.6	30.8
18.0%	702.8	43.2	600.8	28.9
19.0%	662.4	40.4	573.6	27.1
20.0%	624.6	37.8	548.2	25.5
Average		56.0		37.26

The cost of capital over time is exponential, not linear, so the dollar price of risk descends as the discount rate gets higher. This means the cost of being wrong at the upper end of the range (say, 19% to 20%) is cheaper than at the lower end. In every parallel case, however, the cost of being wrong about the Project A discount rate is measurably higher than Project B. For example, being wrong about the Project A 11% discount rate would lead to an immediate loss of $64K if the 12% rate were the correct one. For Project B, a comparable 11% to 12% correction would only equate to a $43K loss. Similarly, if the mistake was misjudging an 11% rate for the correct one of 15%, this would amount to a $234K error for Project A, but only a $155K loss for Project B. In this sense, Project A can be described as being the more risky of the two cash flows – because the unit cost of being wrong is unquestionably higher.

The only time this logic would not prevail would be in those cases where there was a substantial difference in the assessed risk-rates of the two projects. For example, if Project A is believed to have an appropriate risk-adjusted rate of 18%, the cost of being wrong is $40K for the next 100 basis points. If, at the same time, Project B is believed to be correctly risk-compensated at 10%, the cost of being wrong at this level is $46K. This 100 basis point error will be more expensive than with Project A.

OBSERVATIONS AND CONCLUSIONS

- Year-over-year cash flow volatility is an indicator of risk, but the range of dispersion upon which the annual cash expectations are estimated is also a prime determinant of risk
- Each single annual cash estimate is assumed to enjoy the same degree of confidence (i.e. the standard deviations are all proportionately consistent) as every other year
- The FMV is inversely related to the degree of risk in the cash flows
- One of the most effective means of mitigating risk for the purchaser is to negotiate, if possible, a below FMV Purchase Price (or include post-purchase performance concessions)
- The Semi-Variance can be used to identify the proportionate amount of variability in the predicted cash flows above or below a given target return
- The 'cost of being wrong' gives us a present value dollar amount of loss incurred should the discount rate be misjudged. It can be an important comparator of risk when judging between the attributes of two independent investments

We have discussed the relative risk of Project A vs. Project B in detail and yet, still have not referenced one single Market indicator and have no clue as to what the appropriate risk-adjusted returns would be for either project. Indeed, we have no idea even what industry(s) these projects might be in. And yet, with the application of just some very basic statistical analysis we have been able to identify some key risk indicators that should be of interest to our client. While Project A is, undoubtedly, the 'riskier' cash flow in this instance, it does also offer the greater absolute upside potential if the client is a particularly adept negotiator.