Introduction to Risk, Return and the Opportunity Cost of Capital

Jörg Wildförster, 22/09/2010

Investment risk

What is *investment risk* and how is it defined?
How can investment risk be calculated? And finally, what could be done to avoid or minimize the risk for an investment? These three questions could be asked when thinking for the first time about investment risk. Without having studied this topic all three questions could be answered on a gut level. The first question could be answered with; investment risk is the discrepancy between the expected and the real results. Secondly, experiences in the past could be helpful to calculate investment risk in the future. Finally, instead of putting all our eggs in one basket the portfolio of investments should be diversified. Unexpectedly, that is a close representation of how it is done.

Historical review of capital markets and risk premium

A century of capital market history

For securities there are enormous amounts of data available. The focus will be on the performance of three portfolios of U.S. securities measured since 1900. The information is based on a study by Dimson, Marsh and Staunton.

1. A portfolio of treasury bills
   (U.S. government debt securities maturing in less than one year)
2. A portfolio of U.S. government bonds
3. A portfolio of U.S. common stocks

These types of investments offer different degrees of risk. There is no risk of default and relatively stable prices due to short maturity, if looking at treasury bills. It is a safe investment. The only uncertainty is the inflation and the resulting *real* rate of return. The focus is only on the *nominal* rate of return, so inflation does not need to be involved; it should just be kept in mind. Government bonds are assets whose price fluctuates as interest rates vary. Investing in common stocks implies that investors share the entire ups and downs of the issuing companies. The average annual rate of return from 1900 to 2006 was about 4 percent for U.S. treasury bills, 5.2 percent for U.S. government bonds and 11.7 percent for U.S. common stocks. The difference in the average annual rate of return of common stocks and treasury bills (7.6 percent in this case due to rounding) is called (historical) *risk premium*. Please observe that the average annual rates of return of the mentioned U.S. securities are an arithmetic average. In other words the annual returns were added and divided over a period of 107 years.

Someone may ask why the period under consideration is such a long time and why the focus is only on the U.S. market to define our historical risk premium. Firstly, the only significant way of measuring annual rates of return on common stocks is to look at a very extensive period – due to the high fluctuation. Secondly, the study mentioned above leads to the conclusion that the expected risk premium was the same in each industrial nation.

Using historical evidence to evaluate today's cost of capital

In the previous paragraph the average rates of return in the past were examined. Now the question arises how to estimate the expected rates of return respectively the cost of capital today? Let us assume there is an investment project with the same risk as the market portfolio of the Standard and Poor's Composite Index. What rate should be used to discount this project's forecasted cash flows? Obviously this should be the currently expected rate of return on the market portfolio. This will be defined as market return r_m.

Assuming that the future will be comparable to the past and that today's investors expect the same rates of return as in the past; r_m should be set to 11.7 percent, the average of past market returns for U.S. common stocks. As explained above the rates of return are an average of 107 years. Imagine for example that treasury bills were offering 15 percent – as they did in 1981. Would investors then be content to hold common stocks offering an expected return of 11.7 percent? It seems to be unlikely.

That is why r_m is defined as the sum of the current risk-free interest rate (e.g. U.S. treasury bills) and the risk premium.

> Example: In 2006 treasury bills were about 5 percent. Adding on the risk premium 7.6 percent, the expected rate of return was 12.6 percent.
> $r_m(2006) = 5.0\% + 7.6\% = 12.6\%$

The crucial assumption here is that there is a normal, stable risk premium on the market portfolio, so that the expected *future* risk can be measured by the average past risk premium.
Not even with data spanning over 100 years is it possible to estimate the market risk premium, nor is it sure that investors today are expecting the same reward for risk as they did 50 or 100 years ago. All this leaves plenty of room for arguments about what the risk premium *really* is. The reasons why the historical average might overstate today's demanded risk premium are the fact that nowadays the economies and markets seem to be more stable than they were in the past; furthermore, it is easier today for investors to diversify away part of their risk. On the other hand there are probably a few who would expect that today's risk premium is higher than the historical one.

Dividend yields and the risk premium

By having a closer look to the constant-growth model there is another way to calculate the risk premium. If stock prices are expected to keep pace with the growth in dividends, the expected market return is the sum of the dividend yield and the expected dividend growth. The averaged amounts since 1900 in the United States were 4.4 percent for dividend yields and about 5.6 percent for the annual growth in dividends. Assuming that the dividend growth is representative for the investors' expectation, the expected market return over this period was

$$r = \frac{DIV_1}{P_0} + g = 4.4 + 5.6 = 10.0\%$$

respectively 6 percent above the risk-free interest rate. This figure is 1.6 percent lower than the *realized* risk premium of the historical average annual rate of return of U.S. common stocks. Consequently a reduction in the dividend yields would appear to indicate a reduction in the risk premium that investors can expect over the following few years. So, when yields are relatively low, companies may be legitimated in reducing their estimate of required returns over the next few years.

However, changes in the dividend yield give companies no indication about the expected risk premium over the next coming 10 to 20 years. To conclude this debate, it is not possible to estimate what returns investors expect. Many financial economists rely on the evidence of history and work with a risk premium of about 7.5 percent. Others use a somewhat lower figure.

Portfolio risk

The discount rate for safe projects could be assumed by looking at risk-free interest rate (e.g. U.S. treasury bills), while the one for average-risk projects could be estimated by examining the sum of risk-free interest rate and average risk premium. However, it is still undefined how to expect discount rates for assets that do not fit into these simple categories.

Measuring variability

Variability in this sense shows the spread of rates of return in a certain period of time. The variability is measured by *variance* and *standard deviation*. The variance of the market return is defined as

$$\sigma^2_{\tilde{r}_m} = \text{the expected value of} (\tilde{r}_m - r_m)^2$$

where $\tilde{r}_m$ is the actual return and r_m is the expected return. The standard deviation is the square root of the variance

$$\sigma_{\tilde{r}_m} = \sqrt{\sigma^2_{\tilde{r}_m}} \; .$$

The following example illustrates a procedure with which the variability of any portfolio of stocks or bonds could be estimated. The possible results and their probabilities have to be identified and used in the calculation model.

> Example: There is the possibility to play the following game. The initial investment is $100. Then two coins are flipped. If both coins turn up heads, the return would be $140, or plus 20 percent per turn up heads (probability 0.25 percent). If both coins turn up tails the return would be $80, or minus 10 percent per turn up tails (probability 0.25 percent). The third possibility, each coin shows something different; the return would be $110 respectively plus 20 percent for the turn up heads and minus 10 percent for the turn up tails (probability 0.5 percent).
>
> Exp. return = $(.25 \cdot 40) + (.25 \cdot (-20)) + (0.5 \cdot 10) = 10\%$
>
> $\sigma^2 = (40 - 10)^2 \cdot .25 + (-20 - 10)^2 \cdot .25 + (10 - 10)^2 \cdot .5 = 450$
>
> $\sigma = \sqrt{450} \approx 21$

> The games expected return is 10 percent, the variance of the percentage return is 450 and the standard deviation is 21. The standard deviation is in the same unit as the rate of return, meaning that the game's variability is 21 percent.

As we did in the game with the coins, the risk of an asset can also be expressed by all possible outcomes and the probability of each. However, this is theoretical. In daily business it is often impossible. Therefore variance and standard deviation are used for summarizing the spread of possible outcomes. Unfortunately the probability of each outcome is not constant and cannot be found for example in a newspaper. Once more a look to the past could be helpful. The assumption that portfolios with high variability in history also have the least predictable future performance seems to be reasonable. From 1900 to 2006 the annual standard deviation for treasury bills was 2.8, for government bonds it was 8.1, and for common stocks it was 19.8. Please observe that just as today's expected risk premium may be different from the historical one, today's market variability may be different as well.

Diversification reduces risk

By observing ten well-known U.S. common stocks for a five-year period it is surprising how the estimated standard deviations are spread. While the standard deviation of the U.S. market portfolio was about 16 percent the common stock with the highest standard deviation of the analyzed quantity was about 56 percent. That means that this stock was over three times more variable than the market portfolio in the five-year period. This raises the question why the variability of the market portfolio does not reflect the average of its components. The simple but conclusive answer is that *diversification reduces variability*. Diversification works because stocks are not perfectly correlated, meaning prices of different stocks do not move exactly together. Even a little diversification can provide a substantial reduction in variability.

> Example: The standard deviation during the five-year period, ending in June 2006, was about 29.7 percent for IBM and about 29.8 percent for Boeing. Combining the two stocks in equal shares the standard deviation of the portfolio was about 21.5 percent. This is about 8 percent less than holding just one of them.

The enormous improvement of reducing the variability works by combining up to about ten stocks. If the number of securities would be increased beyond, say 20 or 30, the improvement is less significant. The risk that can be potentially eliminated by diversification is called *unique risk*. The basic cause for unique risk comes from the fact that many of the perils that surround an individual company are particular to that company and perhaps its immediate competitors. However, there is also some risk that you cannot avoid regardless of how much you diversify. This risk is called *market risk*. Market risk stems from the fact that there are superior perils that affect all businesses. Due to market risk stocks have a tendency to move together no matter which industrial sector they belong to.

Calculating portfolio risk

Diversification works because stocks are not perfectly correlated. This was the intuitive idea given in the last chapter. Now it should be analyzed more in detail how the risk of a portfolio is related on its components risk. If there is a portfolio consisting of two stocks e.g. Wal-Mart and IBM and these two stocks would move in perfect lockstep the calculation of the expected return respectively the standard deviation is easy.

> Example: 60 percent of the portfolio is invested in Wal-Mart (stock 1) and the remainder is invested in IBM (stock 2). Over the coming year the expected return for Wal-Mart is 10 percent and for IBM it is 15 percent. The expected return r_p on the portfolio is
>
> $r_p = (.6 \cdot 10) + (.4 \cdot 15) = 12\%$.
>
> In the past the standard deviation of IBM was about 29.7 percent and for Wal-Mart about 19.8 percent. Assuming that these figures are also representative for the future the weighted standard deviation of the two stocks is
>
> $\sigma_p = (.6 \cdot 19.8) + (.4 \cdot 29.7) = 23.8\%$.

In reality Wal-Mart and IBM do not correlate perfectly; perfect correlated stocks are unlikely. The exact method for measuring the (two stock) portfolio risk, is to calculate the *portfolio variance* respectively the *portfolio standard deviation* which is simply the square root of the variance.

$$\text{Portfolio variance} = x_1^2\sigma_1^2 + x_2^2\sigma_2^2 + 2(x_1 x_2 \rho_{12} \sigma_1 \sigma_2)$$

x_n denotes the proportion of stock n compared to the portfolio. ρ_{12} denotes the correlation coefficient between stocks 1 and 2.

The product of the correlation coefficient and the two standard deviations is the *covariance* of the two stocks. The covariance, as the word says, expresses the degree to which stocks correlate to each other. If two stocks tend to move together the covariance and the correlation coefficient are positive. They are zero if the stocks are unrelated and negative if they tend to move opposite.

> Example: The correlation coefficient of Wal-Mart and IBM was 0.35 in the past. Expecting the same for the future and using the proportions and standard deviations as in the example before, the portfolio variance is
> $= (.6)^2 \cdot (19.8)^2 + (.4)^2 \cdot (29.7)^2 + 2(.6 \cdot .4 \cdot .35 \cdot 19.8 \cdot 29.7)^2$
> $= 381.1$
> $\sigma = \sqrt{381.1} \approx 19.5\%$.
> The expected standard variation is 19.5 percent. Compared to investing in Wal-Mart only, the portfolio risk is reduced by adding a stock with a higher risk due to diversification.

How individual securities affect portfolio risk

The risk of a security consists of unique risk and market risk. In a portfolio unique risk can be eliminated by diversification, market risk cannot. This fact leads to an important conclusion: *The risk of a well-diversified portfolio depends on the market risk of the securities included in the portfolio.* The sensitivity of a security's market risk to market movements is called *beta* (β). The market portfolio has by definition a beta of 1.0. Stocks with a beta higher than 1.0, tend to amplify the overall movements of the market. A beta between 0 and 1.0 signals a move in the same direction as the market but not as far. Stocks with a high beta have a high market risk and often have a high standard deviation. The risk of a well-diversified portfolio is proportional to the portfolio beta, which equals the average beta of the securities included in the portfolio. Hence, the risk of a well-diversified portfolio with a beta of 0.5 is half of the risk of the market portfolio.
Beta is defined as

$$\beta_i = \frac{\sigma_{im}}{\sigma_m^2}$$

where σ_{im} denotes the covariance between the stock returns and market returns and σ_m^2 is the variance of the returns on the market.

Example: The following table shows the returns over a sixth-month period on the market and on the stock of Anchovy Queen restaurant chain. Both provided an average return of 2 percent. Anchovy Queen's stock was particularly sensitive to market movements, $\beta = 1.5$. In other words, Anchovy Queen is one-and-a-half times as volatile as the market.

Month	Market return	Anchovy Q return	Deviation from average market return	Deviation from average Anchovy Q return	Squared deviation from average market return	Product of deviations from average returns
1	-8%	-11%	-10	-13	100	130
2	4%	8%	2	6	4	12
3	12%	19%	10	17	100	170
4	-6%	-13%	-8	-15	64	120
5	2%	3%	0	1	0	0
6	8%	6%	6	4	36	24
Average	2%	2%		Total	304	456
Variance = σ_m^2 = 304/6 = 50.67						
Covariance = σ_{im} = 456/6 = 76						
Beta (β) = σ_{im}/σ_m^2 = 76/50.67 = 1.5						

Calculating the variance of the market returns and the covariance between the returns on the market and those of Anchovy Queen. Beta is the ratio of the covariance to the variance.

Diversification and value additivity

Diversification reduces risk and therefore makes sense for investors. Does that also mean that firms should do it?
If they do so, each new project has to be analyzed in relation to its value as a potential addition to the firm's portfolio of assets. First of all, a company cannot easily diversify. While an investor could invest in the steel industry this week and pull out next week, a firm cannot do that. The time and expense for a firm to acquire a steel company or to start up a new steel-making operation would be unimaginable. Secondly, as long as investors have the easy option to diversify on their own account, they no longer pay for a firm that offers them diversification. In addition to that, if investors have a sufficiently wide choice of securities, they will not pay any less because they are unable to invest separately in each factory. Consequently in countries with large and competitive capital markets diversification does not add or subtract from the value of a firm. The total value is the sum of its parts. This concept is called *value additivity* and for a two asset firm, asset A and B, it is defined as

$$PV(AB) = PV(A)+PV(B).$$

PV(A) denotes the value for asset A, PV(B) for asset B. The sum is the market value of a firm holding only these two assets.