13.15: Risk and Return Guided Tutorial (CH 7)

The purpose of this guided tutorial is to walk through the process of calculating and interpreting several of the concepts from Chapter Seven on Risk and Return. Please note that this is not a substitute for the chapter in the text, but a complement to it.

We are going to start by estimating the expected return and standard deviation of a single security. This is used when evaluating investments in isolation (sometimes referred to as “stand alone” situations). In order to do these calculations, we are going to start with a probability distribution of possible outcomes. Our probability distributions (referred to as discrete probability distributions because there are a limited number of outcomes) are not designed to capture every possible event that could happen in the real world (as there are potentially infinite outcomes), but instead just to provide a quick overview of some possible scenarios. Once we calculate the expected return and standard deviation, we can evaluate which investment is better by itself.

The next step is to look at the portfolio impacts. While in practice, many people hold portfolios of tens (or hundreds) of different investments, we are going to limit our calculations to a two-stock portfolio to keep things “simple” (the standard deviation formula really blows up as we add more and more securities to the portfolio to the point where it can be very time consuming to do by hand once we get more than a handful of securities). Here we will introduce correlation (how the returns from a pair of securities move together) and consider how it impacts the risk of our portfolio. Portfolios (and the concept of diversification) are a critical element to investing and your personal financial decisions as you manage your real-world finances throughout your life.

Once we’ve gone over the portfolio, we will introduce another measure of risk (beta) that becomes critical once we’ve moved on to larger, well-diversified portfolios. Beta gives us a measure of market risk (which we introduced earlier). When we are evaluating securities in isolation, we want to use standard deviation (total risk ⇒ firm-specific risk PLUS market risk). However, when we are adding stocks to a well-diversified portfolio, we no longer worry about the firm-specific risk (as it has been diversified away) and now must focus on the market risk (beta).

From there, we will introduce the Security Market Line which attempts to integrate the concept of market risk (beta) and required return for a security. In other words, what rate of return do we need in order to compensate us for the risk of investing in the stock? We can then tie that back into expected returns and stock valuation issues.

Please **actively** work through this guided tutorial (doing the calculations and thinking about the issues) instead of just reading through it. You will get a lot more out of it if you do so.

Example One: Big Oil, Inc.

Scenario	Probability	Return
Oil Falls to $35	0.05	-40%
Oil Falls to $55	0.20	-20%
Oil Goes to $80	0.50	12%
Oil Goes to $100	0.25	36%

Calculate the Expected Return and Standard Deviation for Big Oil, Inc. given this probability distribution.

As Big Oil, Inc. is an oil producer/refiner, it stands to reason that their stock price will be partially dependent on the price of oil. As oil increases in value, the stock price should go up as the value of Big Oil’s oil assets increases (and vice-versa). Note that Big Oil is impacted by a lot more than the price of oil (labor costs, regulations, economic growth, etc.) A real-world analysis would be more complex and have several more scenarios. However, for the purposes of this class, we wanted to keep the process relatively straight-forward. Here, this probability distribution is telling us that there is a 5% chance that we will lose 40% of our investment over the next year; a 20% chance that we will lose 20% of our investment over the next year and so on. Also note that for any probability distribution you use to get expected return and standard deviation, your probabilities need to sum to 1.00.

Calculate Expected Return for Big Oil Example

To calculate the expected return, we just take the probability times the return for each possible outcome and then sum them up. Specifically for this problem, we get

$$\bar{k}=\sum_{i=1}^{n} P_i k_i$$
$$\bar{k}=0.05(-40)+0.20(-20)+0.50(12)+0.25(36)$$
$$\bar{k}=-2\%+-4\%+6\%+9\%$$
$$\bar{k}=9\%$$

So, our expected return is 9%. This does not mean we will get 9% as our rate of return. It also does not mean that 9% is our most likely return. As a matter of fact, based on our probability distribution, 9% isn't even a possible outcome and 12% is our most likely return. Instead, what this means is that if you could repeat next year an infinite number of times, sometimes you'd do better than 9% and sometimes you'd do worse, but on average you would have a 9% rate of return.

However, knowing that 9% is your expected return is only part of the story. An analogy that works here is to imagine that you don't know how to swim, but must walk across a lake that is, on average, 3-feet deep. While you may be perfectly fine standing in 3 feet of water, it is pretty easy to see the problem here. The average depth doesn't help you when you get to the spot where the water is 15-feet deep. It is not enough to know the average depth, you need to know the variation in depth. We can get a feel for this through the standard deviation.

While it is not a statistically precise definition, think of standard deviation as a measure of how reliable the expected return is. When you have a low standard deviation, that means that most of the time your actual return will end up close to your expected return. When you have a high standard deviation, that means that you are likely to have an actual return that is significantly higher or lower than your actual return. Note that stock returns are not normally distributed, so don't think of standard deviation as a precise confidence interval. As standard deviation is a measure of how reliable our expected return is, we can think of it as a measure of risk.

Calculate Standard Deviation for Big Oil Example

To calculate the standard deviation, we (1) take the forecasted return for that outcome less the expected return, (2) square that, (3) multiply by the probability of that outcome, (4) sum up for all possible outcomes, and (5) take the square root. Specifically, for this problem we get

$$\sigma=\sqrt{\sum_{i=1}^{n} P_i(k_i-\bar{k})^2}$$
$$\sigma=\sqrt{0.05(-40-9)^2+0.20(-20-9)^2+0.50(12-9)^2+0.25(36-9)^2}$$
$$\sigma=\sqrt{0.05(-49)^2+0.20(-29)^2+0.50(3)^2+0.25(27)^2}$$
$$\sigma=\sqrt{0.05(2401)+0.20(841)+0.50(9)+0.25(729)}$$
$$\sigma=\sqrt{120.05+168.2+4.5+182.25}$$
$$\sigma=\sqrt{475}$$
$$\sigma=21.79\%$$

So our standard deviation is 21.79%. Remember, that the higher this number is, the more likely our actual return is to be significantly higher or lower than our expected return. To see some actual standard deviations for real companies, take a look at the Appendix 1 from Chapter Seven that gives you some comparisons. For example, Pepsi has a standard deviation of just under 12% from Jan. 2013 through Dec. 2017 while Boston Beer's (the maker of Samuel Adams beer) is about 30% over that same time frame.

Let me take a moment to tie standard deviation into our earlier discussion of risk where I stated that the degree of risk is a function of the interaction between (a) the likelihood of an unfavorable event and (b) the degree of the unfavorable event. We can see that in the standard deviation formula. The ($k-\bar{k}$) segment of the formula captures the degree of the unfavorable event – the further from our expected return the greater this value will be. The probability captures the likelihood of the unfavorable event. If one (or both) of these values is small, it will have little impact on the standard deviation. On the other hand if they are both larger, it will lead to a larger standard deviation (more risk).

In order to evaluate where Big Oil ranks in terms of risk and return, we need to compare it to an alternative. So, next let's move on to our next example.

Example Two: Stag Tractors

Scenario	Probability	Return
Recession	0.15	-60%
Slow Economic Growth	0.50	9%
Strong Economic Growth	0.35	50%

Calculate the expected return and standard deviation for Stag Tractors given this probability distribution.

Please try to do these calculations on your own first before moving on.

Calculate Expected Return for Stag Tractors Example

$$\overline{k}=\sum_{i=1}^{n} P_i k_i$$
$$\overline{k}=0.15(-60)+0.50(9)+0.35(50)$$
$$\overline{k}=-9\%+4.5\%+17.5\%$$
$$\overline{k}=13\%$$

Calculate Standard Deviation for Stag Tractors Example

$$\sigma=\sqrt{\sum_{i=1}^{n} P_i (k_i-\overline{k})^2}$$
$$\sigma=\sqrt{0.15(-60-13)^2+0.50(9-13)^2+0.35(50-13)^2}$$
$$\sigma=\sqrt{0.15(-73)^2+0.50(-4)^2+0.35(37)^2}$$
$$\sigma=\sqrt{0.15(5329)+0.50(16)+0.35(1369)}$$
$$\sigma=\sqrt{799.35+8+479.15}$$
$$\sigma=\sqrt{1286.5}$$
$$\sigma=35.87\%$$

Based on our calculations, Stag Tractors has an expected return of 13% and a standard deviation of 35.87%.

Big Oil vs. Stag Tractors

Now, let’s compare these two stocks. Assuming you are a risk-averse investor, which should you choose? The answer is either one, depending on your degree of risk aversion. Big Oil is less risky (lower standard deviation), but it also has a lower expected return. Stag Tractors is riskier, but you are compensated for that additional risk with an extra 4% expected return. Is that additional 4% expected return enough extra compensation for the extra risk? The answer depends on the person. People that are more risk averse (more sensitive to risk) will likely prefer Big Oil. People that are less risk averse will likely prefer Stag Tractors.

What if the expected return for Stag Tractors was only 10%, would some risk-averse investors still choose it? The answer is yes – although the number choosing Stag Tractors would be fewer in this case as fewer people would find the extra compensation (now only 1%) adequate to compensate for the extra risk. As long as the higher risk alternative pays a higher expected return, there should be some people that will choose it. However, the number that choose the higher risk stock will increase as the extra compensation (difference in expected returns) increases.

What if the expected return for Stag Tractors was only 9% (the same as Big Oil) or lower? Now, all risk-averse investors would choose Big Oil as there is no additional compensation for taking on the higher level of risk.

A quick side note. One factor that might influence your decision is your time horizon. Let’s say that I have $10,000 to invest. If I earn a 9% return, I will earn $900 compared to the $1300 I would earn at a 13% return. This is a difference of $400 (keep in mind that is based on expected returns – on average – not actual returns). That $400 difference is probably not very significant when you consider the possible upsides/downsides. However, what if my time horizon was 35 years? If I start with $10,000 and earn 9% for 35 years, I will have $204,140. If I earned 13%, I would have $720,685. Now I end up with about 3.5 times as much wealth. Four percent is not much in one year, but compounded over 35 years it is massive. Therefore, it may make more sense to be willing to take high-risk, high expected-return investments over a longer time frame than if you only have a short time for the return to work in your favor.

Asking you to choose between Big Oil and Stag Tractors (or any two stocks) is really a bit of a false choice. You do not have to choose one or the other…you can choose both. Let’s now move from looking at stocks individually to looking at a two-stock portfolio.

Correlation

Before we can do the analysis of a two-stock portfolio, we need to talk about correlations. Correlations represent the degree to which any two variables move together. In finance, we often focus on the correlation between returns for a pair of securities (note – returns are our focus, not the prices of the securities). Correlations can range from -1.0 (perfect negative correlation) to 1.0 (perfect positive correlation). It is extremely rare for any two variables to have perfect negative, perfect positive, or no correlation (a correlation of exactly zero would be no correlation). Let's look at a few general examples.

Consider if we took a look at two variables ⇒ (1) Hours per week studying and (2) Average test score for each student in the business finance classes over the past few years. We would likely see a positive correlation – but not a perfect positive correlation. Some students study quite a bit and still struggle with the material for a variety of reasons. Other students may not study much, but find the material intuitive and still do well. However, for most students, studying harder will lead to higher test scores. We might expect a correlation probably in the range of 0.6-0.8 (assuming honest answers to hours studied).

Now, let's look at two other variables ⇒ (1) Number of absences during the semester and (2) Average test score for each student in the business finance class over the past few years (physical classes). We would likely see a negative correlation – but not a perfect negative correlation. Some students show up to class, but don't pay attention or just find the material difficult. Some students miss classes, but spend a lot of time outside of class making it up or find the material intuitive. However, in general, the more often people miss class, the lower their test scores tend to be. We might expect the correlation would probably be somewhere around -0.5 to -0.8.

Finally, let's look at two more variables ⇒ (1) Height and (2) Average test score for each student in the business finance class over the past few years. We might expect that there would be no relationship between these two variables (no correlation). However, it is likely that the chance of the correlation being exactly zero would be small. Instead it would likely (by chance) be slightly negative or slightly positive – probably somewhere in the range of -0.10 to 0.10.

For stocks, correlations tend to be small positive values (between 0.10 and 0.50) due to general economic conditions (market risk factors) that impact all stocks. However, firm-specific risk factors keep the correlations from getting too high as they impact each stock independently. Typically, firms in the same industry tend to have higher correlations. Again, refer to the data in Appendix 1 which captures some sample correlations based on monthly returns over the 2013-2017 time frame. A few quick examples would be Pepsi-Coke (same industry) with a correlation 0.69 or Caterpillar-Molson Coors (very different industries) with a correlation of 0.03. Note that in this sample, all but nine (out of 66) of the paired stock correlations are positive and average 0.20. The Aggregate Bond ETF is a portfolio of bonds, not an actual company. This allows us to see how stocks are correlated to the overall bond market. The S&P 500 is a measure of the overall stock market (500 individual companies in one large portfolio) and also not a company. Both the Aggregate Bond and S&P 500 are represented by Exchange Traded Funds (ETFs) which are a type of mutual fund. These allow individuals to own a broad portfolio of securities as if they were a single stock and can be important investment tools. We will come back to the S&P 500 later in this chapter as well.

Now that we understand correlation, we can move to the expected return and standard deviation for a two-stock portfolio.

Example Three: 2-Stock Portfolio

The formula for expected return of a two-stock portfolio is:

$$\bar{k}_p = w_1\bar{k}_1 + w_2\bar{k}_2$$

Where the w stands for weight (proportion of our portfolio in that stock) and $\bar{k}$ stands for the expected return. In our example, Big Oil is stock 1 and Stag Tractors is stock 2. Since we are investing \$4000 into Big Oil and \$6000 into Stag Tractors, our total portfolio value is \$10,000. That means the proportion of our portfolio into Big Oil is 0.4 (\$4000/\$10,000) and the proportion of our portfolio in Stag Tractors is 0.6 (\$6000/\$10,000). Note that our weights should be expressed as decimals and should add to 1.0. Remember from Examples One and Two that the expected return for Big Oil was 9% and the expected Return for Stag Tractors was 13% (I recommend writing these as 9 and 13 in the formula instead of 0.09 and 0.13, but you can take either approach as long as you are consistent – don't mix and match). This gives us

$$\bar{k}_p = (0.4)(9\%) + (0.6)(13\%)$$
$$\bar{k}_p = 3.6\% + 7.8\%$$
$$\bar{k}_p = 11.4\%$$

So, the expected return for our two-stock portfolio is 11.4%.

The standard deviation is a bit more complicated. The formula for this

$$\sigma_p=\sqrt{W_1^2\sigma_1^2+W_2^2\sigma_2^2+2W_1W_2\sigma_1\sigma_2 corr_{1,2}}$$

Our weights are still 0.4 for Big Oil and 0.6 for Stag Tractors. Remember from Examples One and Two that the standard deviation for Big Oil is 21.79% and the standard deviation for Stag Tractors is 35.87%. Like the returns, you can either express these as 21.79 and 35.87 or as 0.2179 and 0.3587 in your calculations. I recommend 21.79 and 35.87, but it is up to you…just don't mix and match. Again, make sure the weights are expressed as decimals and add to 1. The correlation given here is -0.5. So, when we plug these into the formula we get

$$\sigma_p=\sqrt{(0.4)^2(21.79)^2+(0.6)^2(35.87)^2+2(0.4)(0.6)(21.79)(35.87)(-0.5)}$$
$$\sigma_p=\sqrt{75.97+463.20+(-187.59)}$$
$$\sigma_p=\sqrt{351.58}$$
$$\sigma_p=18.75\%$$

Note that in this example, the standard deviation of the two-stock portfolio is actually less than the standard deviation of either stock individually. That is due to the extremely low (actually negative) correlation which means that these stocks have a tendency to move in opposite directions. When Big Oil is doing poorly, Stag Tractors tends to do better. When Stag Tractors is doing poorly, Big Oil tends to do better. While this is a bit of an extreme example (remember, negative correlations are rare), we can diversify (reduce) our overall risk any time the correlation is less than 1.0 (which is almost always going to be the case). By reducing our risk, we mean that the risk of the portfolio will be less than the weighted average of the individual stocks. Two clichés come to mind. The first – "there is no free lunch" – is made false by the second – "don't put all your eggs in one basket". Because diversification does not lower our expected return (it is a simple weighted average that is not affected by the correlation), our "free lunch" is that by holding stocks in a portfolio, we can reduce our risk without lowering our expected return. While we demonstrated this with a two-stock portfolio (to keep the calculations manageable), the more stocks we add to the portfolio (up to a point), the more benefits we get from diversification.

Because of the idea of diversification, we can virtually eliminate the impact of firm-specific risk by holding a large portfolio of 50+ stocks from a variety of different industries (and even countries). However, no matter how many stocks we own we still are faced with market risk. While all stocks are subject to market risk, some stocks are more sensitive to market risk than others. If we have a portfolio of all high market risk stocks, we will have diversified away most of our firm-specific risk, but still have a high risk (and high expected return) portfolio. Alternatively, if we have a portfolio of all low market risk stocks, we will have diversified away most of our firm-specific risk, but now will have a lower risk (and lower expected return) portfolio. Since we can eliminate most of our firm-specific risk it becomes less relevant. However, since we can not eliminate market risk, we need to be able to measure it. It also should be the primary factor that systematically drives returns. We will measure market risk with a concept referred to as beta.

Example Four: Beta

Next, you estimate the expected return and standard deviation for the S&P 500 (The S&P 500 is a portfolio of 500 large-company stocks) and find that it has an expected return of 10% and a standard deviation of 14%. Big Oil has a correlation with the S&P 500 of 0.75 while Stag Tractors has a correlation with the market of 0.32. Calculate the beta of each stock.

The formula for beta is

$$\beta=\frac{(\sigma_{stock})(corr_{stock,market})}{\sigma_{market}}$$

Looking at the formula, you see the word "market". Theoretically, "the market" refers to a value-weighted portfolio of all investible assets that one can buy. This would include all stocks across the globe, all bonds across the globe, all real estate, precious metals, precious gems, commodities, collectibles (paintings, antiques, etc.) and anything else that someone could purchase as an investment. From a practical perspective, "the market" is too broad to measure. Therefore, we often see people use a simple approximation such as the S&P 500. The S&P 500 is a "market portfolio" of 500 larger firms grouped together as one large portfolio (with the largest companies making up a larger portion of the portfolio). On the handout with the sample correlations and standard deviations, I've included the S&P 500 as a measure of the market. Here, all the necessary values are provided in the example. This allows us to calculate the beta of Big Oil and Stag Tractors as follows (remember that the standard deviation of Big Oil was 21.79% and the standard deviation of Stag Tractors was 35.87%):

$$\beta_{BigOil}=\frac{(21.79)(0.75)}{14}=1.17$$

$$\beta_{StagTractors}=\frac{(35.87)(0.32)}{14}=0.82$$

Note that we do not use the expected return on the market here (we will come back to that in the next example). Also, note that while we previously stated that Stag Tractors was riskier than Big Oil when we were evaluating with standard deviation, beta tells us a different story. The reason that standard deviation can give us a different result than beta is that they are measuring two different types of risk. Standard deviation measured total risk (firm-specific risk plus market risk). This was important when evaluating investments on their own as the firm-specific risk had not been diversified away. Beta measures market risk, which is critical when evaluating stocks as part of a well-diversified portfolio because then firm-specific risk is no longer relevant. While Stag Tractors is riskier by itself, its lower correlation to the market (again, remember these are our hypothetical example numbers from these practice problems) allows it to have less market risk and be less risky as part of a well-diversified portfolio.

The average risk level for market risk is 1.0. The further the beta falls below 1.0, the lower the market risk. The higher the beta rises above 1.0, the greater the market risk. The data from Appendix 1 gives some sample betas for the time frame of 2013-2017. For example, the beta for Boston Beer is a low 0.51 indicating very low market risk. This makes sense as we would not expect Boston Beer's cash flows to be very sensitive to overall economic factors – indicating lower market risk. Alternatively, Amazon has a high beta of 1.47. Again, this makes sense as Amazon (an online retailer and technology company) will see it's returns be very sensitive to the overall market – indicating higher market risk. While beta is estimated off of past data, what we really need is the beta for the upcoming period. Given that beta changes over time, we need to think of beta as a rough approximation, not an exact value.

Because investors know that they can simply diversify away firm-specific risk, the critical risk that they want to focus on (and get compensated for) should be market risk – measured by beta. Since market risk drives investors required returns (higher betas should have higher required returns as they are riskier) we need a way to formalize this process. Fortunately, this was developed in the 1960's within a theory referred to as the Capital Asset Pricing Model (CAPM). The CAPM is a broader theory relating risk and return, but the practical application component of the CAPM is the Security Market Line (SML). While technically they are not the same thing (the SML is a part of the CAPM theory), many people use the terms interchangeably. The next example walks us through the SML and its implications.

Example Five: Security Market Line and Required Return

Then, assuming the current risk-free rate of interest is 4.7%, calculate the required return of each stock.

Based on this, should you buy either (or both) of these stocks as part of a well-diversified portfolio?

The formula for the SML is as follows:

$$k=k_{RF}+\beta(\bar{k}_m-k_{RF})$$

Where k represents the required return for the specific stock (or portfolio) in question, k_{RF} represents the risk-free rate of return, β refers to beta, and $\bar{k}_m$ refers to the expected return on the market. Oftentimes, people will refer to the portion $(\bar{k}_m-k_{RF})$ as the market risk premium as it reflects the additional compensation (beyond the risk-free rate) that investors need to compensate them for the risk of owning stocks. The idea is that we know stocks are risky investments and investors don't like risk. Therefore, investors need additional compensation (beyond a risk-free rate) to make it worthwhile to own stocks. However, different stocks have different levels of (market) risk – measured by beta. Higher risk stocks (those with higher betas) should get more than average additional compensation. Alternatively, lower risk stocks (those with lower betas) should get less than average additional compensation. Therefore, we multiply the market risk premium by beta to adjust for the risk of that specific stock and get the proper amount of additional compensation for that specific stock. Then, we add back in the risk-free rate to get the total required return for that specific stock.

If we apply it to our specific stocks, we get the following (remember from example four that the beta for Big Oil was 1.17, the beta for Stag Tractors was 0.82 and the expected return for the S&P 500 was 10%):

$$k_{BigOil}=4.7\%+1.17(10\%-4.7\%)$$
$$k_{BigOil}=4.7\%+1.17(5.3\%)$$
$$k_{BigOil}=4.7\%+6.20\%$$
$$k_{BigOil}=10.90\%$$

$$k_{StagTractors} = 4.7\% + 0.82(10\% - 4.7\%)$$
$$k_{StagTractors} = 4.7\% + 0.82(5.3\%)$$
$$k_{StagTractors} = 4.7\% + 4.35\%$$
$$k_{StagTractors} = 9.05\%$$

Thus, our required return for Big Oil is 10.90% and our required return for Stag Tractors is 9.05%. Remember from Examples One and Two that the expected return for Big Oil was 9% and the expected return for Stag Tractors was 13%. For Big Oil, the expected return (9%) is less than the required return (10.90%) meaning that we will not (on average) earn enough to compensate us for the risk. Therefore, we should not buy Big Oil (assuming we don't already own it) or we should sell Big Oil (assuming we do already own it). Alternatively, the expected return for Deere (13%) is more than the required return (9.05%) meaning that we will (on average) earn more than enough to compensate us for the risk. Therefore, we should buy Deere stock. We could also use these required returns in the stock valuation models introduced in Chapter Five (remember that back then, the required return was given).

There are a couple of challenges in using the Security Market Line. First, there is no truly risk-free rate. All investments have some level of risk. However, typically the yield on Treasury securities (debt issued by the US Federal Government) is used as the risk-free rate. There is some debate about whether short-term debt (3-month Treasury bills) should be used (as shorter-term securities are less sensitive to changes in the market rate of interest and thus less risky) or long-term (10-year Treasury bonds) should be used (as stocks are a long-term investment and using 3-month securities introduces reinvestment risk as we don't know what rate will be in place after the first 3 months has gone by). I prefer the 10-year Treasury. In addition, there is some debate about the appropriateness of using Treasury debt in today's economic environment as the US no longer has a AAA rating from S&P and the increasing levels of US Federal Debt may introduce some default risk. However, at least for now, investors are accepting incredibly low yields on US Treasury debt which indicates that, while it may not be 100% risk-free, it is still viewed as a very low risk investment.

A second problem was referred to earlier which is that ideally we need the beta for the future time horizon and we can only estimate the past. As this estimation risk means we do not know the true beta, our SML results will be subject to estimation error as well.

A third problem is that the market risk premium is unknown as there is no way to determine exactly what the expected return on the market is. We know that the market risk premium needs to be positive (risk-averse investors would not take on the risk of holding stocks if they expected to earn less than the risk-free rate), but we don't know whether it should be 3%, 5%, 7% or some other value. Historically, it has been estimated to range from about 3% to 7% depending on economic conditions and recent performance of the stock market. When the economy is strong and the recent performance is strong, investors are more willing to hold stocks and the market risk premium shrinks. When the economy is weak and the recent performance of stocks has been strong, investors are hesitant to hold stocks and need more compensation to get them to do so – making the market risk premium larger.

There are two important implications of the Security Market Line (assuming it is valid):

1. The only factor that should systematically differentiate returns across portfolios is beta. This is because the only factor that is different from one stock to the next is beta (both the risk-free rate and the expected return on the market are not dependent on what stock we are looking at).
2. High beta stocks should (on average) earn higher returns over time while low beta stocks should (on average) earn lower returns over time due to the risk differential.

Note that this does not say firm specific factors (company X gets hit with a lawsuit or company Y develops a great new product) don't influence stock returns, just that they aren't systematic. Instead, they will cancel out in larger portfolios and become irrelevant. It also doesn't say every high beta stock will outperform every low beta stock (again, firm specific issues can impact individual stocks) or that high beta stocks as a group will always outperform low beta stocks (actually high beta stocks should do worse in periods where the market as a whole does poorly).

Unfortunately, studies indicate that neither of these implications holds consistently. First, other factors (listed below) have been shown to impact returns on a systematic basis. Second, it is not clear that high beta stocks do (on average) earn higher returns over time.

Other factors that have been shown to influence stock prices include:

1. Firm Size – small firms tend to earn higher returns than large firms even after controlling for beta
2. MV/BV ratio – "value" firms (firms with low MV/BV ratios) tend to earn higher returns than "glamour" or "growth" firms (firms with high MV/BV ratios) even after controlling for beta

3. Momentum – stocks that have outperformed the market over the last six months are likely to earn higher returns over the next six months than stocks that have underperformed the market over the last six months, even after controlling for beta.

These issues have led to the development of alternative asset pricing models (which go beyond the scope of this course).
