ECONOMICS 134 (NEHRING)

HANDOUT ON THE "SHORT-RUN MACRO-EQUILIBRIUM MODEL" OF THE EQUITY RISK PREMIUM.

Introduction

The old investment mantra tells us that "if you bear more risk, you get a higher return". In this section we explore the reasoning behind this rule, and a mathematical model of portfolio selection demonstrating this 'fact'.

To reinforce what has been gone over in class, let's lay down the behavioral assumption underlying portfolio selection:

1) Agents always prefer a higher expected return to a lower expected return. This makes intuitive sense: people always want more than less of a 'good' thing, in this case money.

2) Agents are risk averse, i.e. always prefer less risk to more. In the case of investment risk, we are talking about the variability of returns. Certain investments are "high risk" because their returns have a relatively high probability of large surprises, i.e. extreme deviations from the mean. This assumption is a bit stronger than the previous one. The objection might be raised that it is possible for people to be risk neutral, and maybe even risk loving. While rare, both of these ideas can be incorporated into our model, but for the most part we will work with the assumption of risk aversion. These two assumptions tell us what agents' indifference curves will look like: they will show a trade off between return and risk.

The Model

The basic model starts with the two assumptions outlined above. Next we introduce two assets: the risk-free short-term bond, F, and the risky market portfolio (stocks), M. The bond is risk-free, because in this model risk is defined as variability of return over the short-term return. The stocks are risky for two reasons: one, because it is not generally know how much dividends will be paid out, and, two, stock prices fluctuate making capital gains uncertain. The choice our agents must make is to determine what percentage of their investments to put in the risk-free bond, and what percentage to put in the risky stocks. (Note: Although there are many stocks in reality, here there is only one "stock" to choose from for now. Think of it as investing in a broad based mutual fund.

We denote it as M because it is like investing in the 'Market'.)

Mathematically, we denote return as μ. μ denotes the 'expected return' of an investment or portfolio. We denote σ as risk. Recall, risk is the variability of returns. In these models, we use standard deviation as our measure of risk. Figure #1 shows what indifference curves will look like in risk-return space with the different assumptions of risk aversion, neutrality and loving. This is the agent side of the problem. Figure #2a shows the location of the two assets in risk-return space. This is the basis of the market side of the problem. If we invested in only bonds, we would have a portfolio corresponding to point F. If we invested only in stocks, we would have a portfolio corresponding to point M.

What if we invested a proportion of assets, w, in stocks, and the rest, $1-w$ in bonds? Call this portfolio P_w. Where in the space would our portfolio, P_w, lie? First, we can get the vertical coordinate by calculating the expected return of P_w. Second, we can get the horizontal coordinate by first calculating the variance of P, and then by taking the square root to get the standard deviation:

Expected value: $E[P_w] = E[w \cdot M + (1-w) \cdot F] = w \cdot E[M] + (1-w) \cdot E[F] = w \cdot \mu_M + (1-w) \cdot \mu_F$

Variance: $V[P_w] = V[w \cdot M + (1-w) \cdot F] = w^2 \cdot V[M] + (1-w)^2 \cdot V[F] + 2 \cdot Cov[M,F] = w^2 \cdot \sigma_M^2 + (1-w)^2 \cdot \sigma_F^2$

$= w^2 \cdot \sigma_M^2$

Standard deviation: $STD[P_w] = \sqrt{V[P_w]} = \sqrt{w^2 \cdot \sigma_M^2} = w \cdot \sigma_M$

Now, notice that this portfolio is on the line between F and M. If you, say, invested half in stocks and half in bonds, you would have a return half way between F and M, and a variance half way between F and M. See figure #2b. If we think of the market as presenting us with a budget set of risk-return combinations we may choose from, we see now that we can expand our idea of possible portfolios to include not only F and M, but the line segment that connects them. In other words, when the agents in our model want to find the portfolio that is best for their preferences, they can choose a risk-return combination that is on the line segment between F and M. This is know as the capital market line, CML, and it represents all the portfolios that an agent can possibly buy. The capital market line can be extended beyond point M, by "borrowing on the margin". This is where agents borrow money, at the risk-free rate, to put into stocks. This would correspond to points where the weight, w, put on stocks, M, is greater than 1, which corresponds to the weight on bonds $1-w$ being negative. This is depicted in figure #2c

The problem is now well defined. The agents in our model want to maximize their utility from within the realm of possible choices, the CML. A typical solution for a risk averse agent is shown in figure #3a. Other solutions are shown in figures #3b and c. In figure #3b, the agent has a strong preference for stocks. In this case, though, there is no borrowing on the margin, so they are constrained to pick a weight on stocks of 1. In #3c, there is borrowing on the margin, so the agent is free to choose a weight of greater than 1.

A Closer Look At Preferences

We now give a more mathematical look at agents' preferences. One way to describe agents' preferences is in terms of the following class of utility functions:

$U(\mu, \sigma) = \mu - \frac{\alpha}{2}\sigma^2$

As above, μ is the expected return of the agent's portfolio, σ is the standard deviation, and σ^2 is the variance. α is known as the "coefficient of risk aversion". Let's examine the meaning of this equation. First of all, an agent's utility is a function of risk and return. As we would expect, if the expected return μ increases, utility will increase. If risk, or standard deviation of return, σ, increases, then utility will decrease. If α, the coefficient of risk aversion, is large then a given increase in risk will have larger negative impact on utility. Thus, we say that α is a measure of how risk averse an agent is.

Notice that this mathematical characterization of agents' preferences can represent risk neutral and risk loving agents as well as risk averse agents. If $\alpha = 0$, then the agent's utility will be $U(\mu, \sigma) = \mu$. In other works, an agents utility will depend only on expected return. This is the case of risk neutrality. If α is negative, then increases in risk will increase utility (the two negatives will cancel each other). This is the case of risk loving.

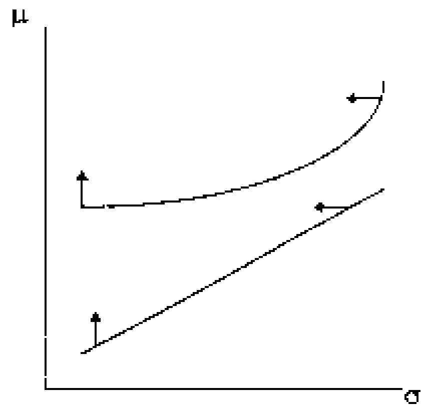

Figure 1a: Risk Adverse Preferences.

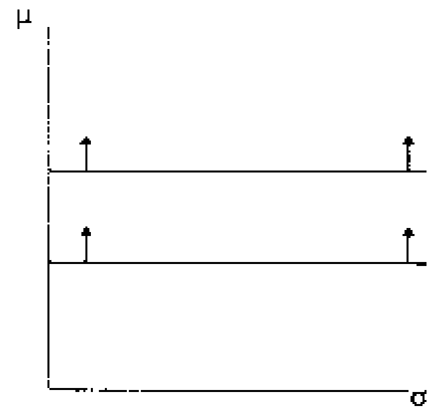

Figure 1b: Risk Neutral Preferences.

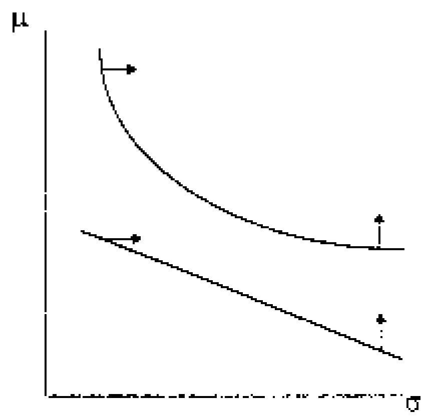

Figure 1c: Risk loving preferences

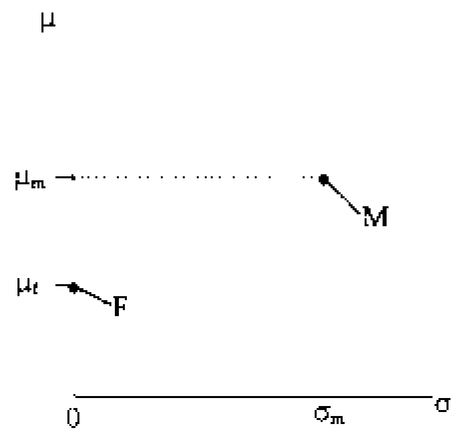

Figure 2a: The riskless bond and risky stock in Risk-Return space

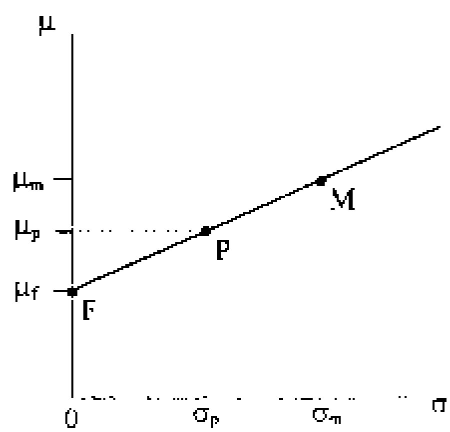

Figure 2b: The portfolio, P, as part stocks and part bonds, in Risk-Return space.

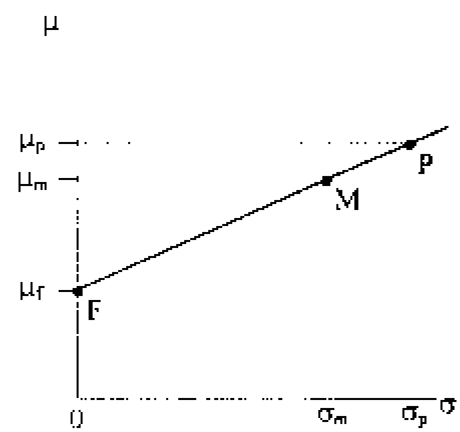

Figure 2c: The portfolio, P, as greater than
100% stocks, financed through borrowing

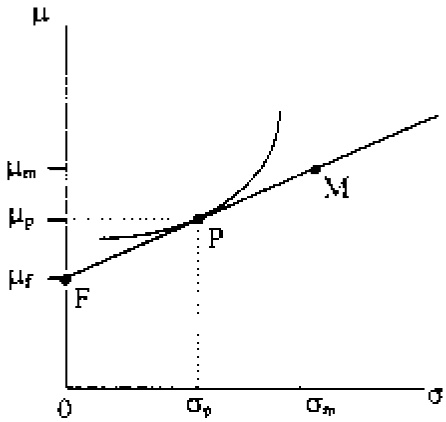

Figure 3a: The portfolio, P, as part stocks and part bonds, in is the portfolio most preferred by this risk adverse investor

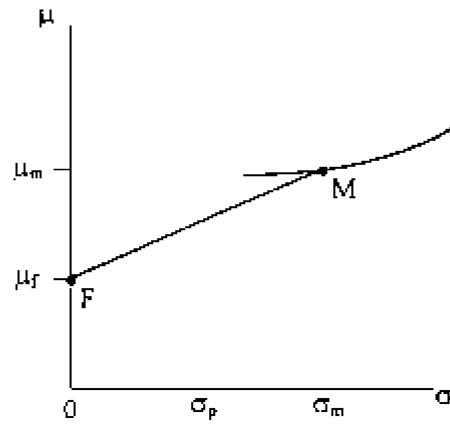

Figure 3b: The portfolio, P, as part stocks and part bonds, in Risk-Return space.

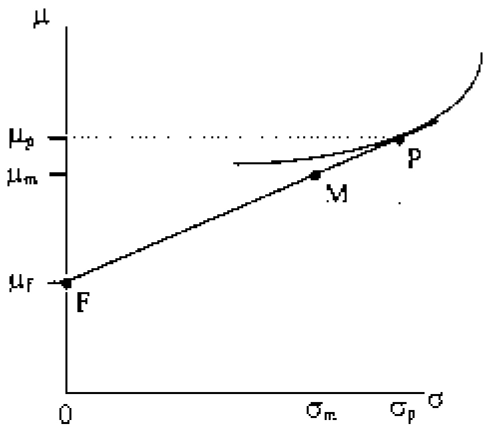

Figure 3c: A person with these risk adverse preferences prefers the portfolio, P, which has greater than 100% stocks, financed through borrowing at the riskless rate.

Application: Short-Run Macro Equilibrium

An important application of our model is to the 'whole' market. The idea here is this: In the market as a whole, there is a fixed amount of stocks. Given this fixed amount of stock, there is a fixed amount of risk associated with it. This risk will not change if the price of the stock changes. So, what we want to do is find the right 'price' for all the risk that is out there. When we say "right price", we mean the equilibrium price for risk, where the aggregate supply equals the aggregate demand. The question is, how do we do this with our model?

First, we need an idea of *aggregate demand*. The demand comes from individuals who buy stocks and bonds. Here, we aggregate them into on "*representative agent*". We imagine the demand side of the market can be represented by one "person", or in the case of our model, one "utility function". This is similar to the case where many individuals have individual demands for a good, say apples. Then, when we look at the market as a whole, we think of there being one aggregate demand curve for apples. So, in our model, we will model aggregate demand as one big agent, and we will call this agent's demand for risk the market's demand for risk.

Second, we need an idea of '*price*'. For 'price' we use the required rate of 'return' of stocks μ_M as a measure of the reward for bearing risk, or equivalently the equity risk premium $\mu_M - \mu_F$. Notice that we could use price instead of return directly, at the cost of complicating the analysis. A stock is a right to a future stream of payments (dividends), and if the price of this 'stream of payments' goes down, the return from investing in it goes up.

Third, we need to introduce an *aggregate supply of risk*. We will look at equilibrium in the short run; since the price-adjustment in capital markets is extremely fast, the short-run can be assumed to be a day, or even less. Thus, it is reasonable to assume that the short-run supply of assets is *perfectly inelastic.*

In this model, the agent demands risks by choosing an optimal portfolio-mix of M and F. Thus, the appropriate notion of the short-run supply of risk is given by the risk σ_T of the portfolio T of all outstanding assets. If the weight of stocks M in the total portfolio T is given by w^*, then the short-run supply of risk is given by $\sigma_T = w^* \cdot \sigma_M$; it is perfectly inelastic, since it does not depend on the return on stocks.

Finally, we need to specify what we mean by equilibrium. In the short run, it is reasonable to assume the risk of stocks σ_M and the risk-free interest rate μ_F as fixed. Asset markets are cleared by an appropriate adjustment of the return on stocks (or, equivalently, of the price of stocks). Intuitively speaking, *the equilibrium rate of return must be such that the representative investor wants to hold exactly the portfolio of existing assets T.*

Consider a given (non-equilibrium) return on stocks; this generates a Capital Market Line (CML). Given the CML, the representative agent's utility curve will be tangent to it at a certain point. This point will mark how much risk the agent (who represents the market as a whole) demands. If the agent demands more or less risk then the market has to offer, then we are out of equilibrium. Let's suppose the agent demands less risk than the market has. This would mean that the price of stocks (i.e. risk) is too high, or in other words the return on stocks is too low. Thus, the return must increase to get the agent to demand more stocks (i.e. more risk).

Now picture this graphically. When the return on stocks increases, the point M moves up vertically, keeping σ_M constant, while increasing μ_M. So, the whole CML slopes up more, and the tangency point will move further up the CML, closer to M. Equilibrium will be reached when the agent demands T, the total portfolio of existing assets. This is where the supply of risk, σ_T, equals the demand for risk as given by the agent's preferences.

Now let's analyze this more mathematically. The point where the agent's indifference curve is tangent to the CML line gives us how much risk she demands. To simplify things mathematically, we say that the agent must pick a portfolio on the CML line. Since the CML line gives us a formula for return in terms of risk, we can plug this into the agents utility function. We can do this because the agent must pick a portfolio on the CML. Sounds confusing? Well, look at these equations, and then re-read this paragraph to see if it makes sense.

The equation of the CML line is: $\mu(\sigma) = (\frac{\mu_M - \mu_F}{\sigma_M}) * \sigma + \mu_F$. Let $\pi = \frac{\mu_M - \mu_F}{\sigma_M}$. We now plug this into the agent's utility function:

$$
\begin{aligned}
u(\mu, \sigma) &= \mu - \frac{\alpha}{2} * \sigma^2 \\
u(\mu(\sigma), \sigma) &= [(\frac{\mu_M - \mu_F}{\sigma_M}) * \sigma + \mu_F] - \frac{\alpha}{2} * \sigma^2 \\
u(\mu(\sigma), \sigma) &= \mu_F + \pi\sigma - \frac{\alpha}{2}\sigma^2
\end{aligned}
$$

We now need to find the σ that maximizes the agent's utility. We do this by taking the first derivative and setting it equal to zero. This gives us:

$$\pi - \alpha\sigma = 0$$

And, solving for σ:

$$\sigma_d = \frac{\pi}{\alpha}$$

This is the agent's demand for risk!

Now, we set supply equal to demand, and see which μ_m will equate them:

$$\begin{aligned}
\sigma_d &= \frac{\pi}{\alpha} = \frac{\frac{\mu_M - \mu_F}{\sigma_M}}{\alpha} = \sigma_T, \text{ which is equal to } w^* \sigma_M. \\
\text{Hence } \mu_M^{eq} &= \alpha w^* \sigma_M^2 + \mu_F
\end{aligned}$$

And so, this is how we get the equilibrium return on risk, μ_M^{eq}.

This makes it possible to perform a variety of comparative statics exercises. For example, if the supply of stocks increases relative to that of bonds (i.e. if w^* increases), one would expect the price of stocks to go down. This is exactly what the above equilibrium formula delivers, since it shows that there expected return μ_M^{eq} will rise in this case.