

Final Exam

ECON 337901 - Financial Economics
Boston College, Department of Economics

Peter Ireland
Spring 2021

Due Saturday, May 15

This exam has three questions on seven pages; before you begin, please check to make sure that your copy has all three questions and all seven pages. Each question will be weighed equally in determining your overall exam score.

To answer questions 1 and 3, you will need to download two files from the course webpage. One, containing monthly data on stock returns from 2016 through 2020, is available either as an Excel spreadsheet or a plain text file at

http://irelandp.com/econ3379/exams/final21s_monthly.xlsx

or

http://irelandp.com/econ3379/exams/final21s_monthly.txt

The other, containing annual data on stock returns from 1927 through 2020, is available in the same two formats at

http://irelandp.com/econ3379/exams/final21s_annual.xlsx

or

http://irelandp.com/econ3379/exams/final21s_annual.txt

Note that Excel spreadsheets can also be read by Google Sheets, so if you don't have Excel on your computer or if you are going to do the calculations in Google Sheets anyway (instructions for this are given on the pages that follow), you can use Google Sheets instead.

Please circle your final answer to each part of each question after you write it down, so that I can find it more easily. If you show the steps that led you to your results, however, I can award partial credit for the correct approach even if your final answers are slightly off.

This is an open-book exam, meaning that it is fine for you to consult your notes and all material from the course and Canvas webpages when working on your answers to the questions. I expect you to work independently on the exam, however, without discussing the questions or answers with anyone else, in person or electronically, inside or outside of the class; the answers you submit must be yours and yours alone.

1. Estimating CAPM Betas

Financial economists and financial market participants typically estimate CAPM betas for individual stocks using 5 years of monthly data. The 5-year data sample is long enough to provide statistically precise estimates. But, by going back only 5 years, these analysts avoid using data from the more distant past that may not be relevant for predicting stock returns in the future.

The monthly data set mentioned above contains data on monthly returns, running from January 2016 through December 2020 on the market as a whole and on six individual stocks: Apple Computer, Ford Motor Company, Google (now officially known as Alphabet Inc), Kellogg, US Steel, and Walmart.

Use these data to estimate the CAPM betas for each of the six individual stocks.

There are at least three ways of doing this. Choose whichever is easiest or most convenient for you.

One way is to load the data into a statistics or econometrics software package like Stata or R, and regress each individual stock return on a constant and the market return. The slope coefficient from this regression equals the CAPM beta: the covariance between the dependent variable (the individual stock return) and the independent variable (the market return) divided by the variance of the the independent variable (the market return).

It is also possible to estimate the slope coefficient from the regression described above using Google Sheets. To do this, download the Excel spreadsheet and open it, instead, with Google Sheets. Then, to estimate the beta for Apple Computer's stock, use the command

```
=slope(C12:C71, B12:B71)
```

For the other five stocks, keep the range of the independent variable B12:B71 as above, but change the range of the dependent variable to match the columns containing data for the other stock returns.

A third way is to use Excel, Google Sheets, or some other computer program to calculate the covariance between the individual stock return and the market return and the variance of the market return, then to divide the covariance by the variance.

Now you know how to estimate CAPM betas!

2. CAPM Betas and Expected Returns

Use the CAPM betas that you estimated above to determine the expected return on each of the six individual stocks.

To do this, you also need to estimate the risk-free rate and the expected return on the market portfolio. For this task, it is useful to find data on a slightly longer sample of data, since interest rates have been unusually low and stock returns unusually high over the past five years alone.

The annual data set mentioned above contains data on annual returns, running all the way back to 1927 and extending through 2020. Because this period includes the Great Depression of 1929 through 1933 and the era of very high inflation during the 1970s, however, it may be too long to reliably predict stock returns in the near future.

Instead, let's use the annual data but focus on just the 21 years from 2000 through 2020. Over that period, the average return on the stock market has been 8.80 percent and the average risk-free rate has been 1.59 percent. These averages suggest setting $E(\tilde{r}_M) = 8.80$ and $r_f = 1.59$ in the CAPM formula

$$E(\tilde{r}_j) = r_f + \beta_j[E(\tilde{r}_M) - r_f].$$

Plugging these numbers in on the right-hand side together with your estimates of the betas from question 1 will produce the CAPM estimates of the expected return on each of the six stocks.

3. Value Stocks and the CAPM

Many successful investors have learned that it is possible to “beat the market,” that is, earn average returns above those provided by the market as a whole, by investing in “value stocks,” that is, shares in old-fashioned companies that most other investors ignore and whose prices are beaten down and by avoiding “growth stocks,” that is, shares in newer and more popular companies whose prices have already been bid up. Warren Buffett, most famously, made his fortune partly through “value investing,” but many others have profited from the strategy as well.

The superior performance of value stocks over growth stocks can be quantified using the “HML” measure developed by Eugene Fama and Kenneth French and described in more detail below. According to this measure, value stocks provided annual returns that were, on average, 5.9 percentage points higher over a long period extending from 1927 through 2006. Interestingly, growth stocks have done much better recently: from 2007 through 2020, they have returned, on average, 6.9 percentage points more than value stocks. Even after accounting for this recent reversal, however, over the full period from 1927 through 2020, the annual return on value stocks remains 4.0 percentage points higher, on average, than the return on growth stocks.

Financial economists will always be quick to point out that higher average or expected returns on value compared to growth stocks do not necessarily violate the implications of the CAPM. Suppose, in particular, that value stocks have higher CAPM betas than growth stocks. Then the CAPM predicts that value stocks will have higher expected returns as well. The interpretation of Warren Buffett’s success would then be: he earned higher returns, but only because he was willing to take on more aggregate risk.

It turns out, though, that two famous research articles, one by Barr Rosenberg, Kenneth Reid, and Ronald Lanstein (“Persuasive Evidence of Market Inefficiency,” *Journal of Portfolio Management*, Spring 1985) and the other by Eugene Fama and Kenneth French (“Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics*, February 1993), find that higher average returns on value stocks compared to growth *do* violate the CAPM. That is, value stocks offer higher average returns than growth stocks, even after accounting for differences in their CAPM betas.

By answering this question, you will see whether or not the earlier results presented by Rosenberg, Reid, and Lanstein and by Fama and French continue to hold in the longer set of annual data, running from 1927 through 2020, mentioned above. This data set runs well beyond those used in the original studies. This exercise is both interesting and important because, as noted above, value investing strategies have performed quite poorly, in recent years, while growth stocks like Apple, Facebook, and Google have provided far superior returns. If you include this recent data in your sample, will it still appear that value stocks outperform growth stocks on average? Are the differences in expected returns still larger than what is predicted by the CAPM? And are the deviations from the CAPM’s predictions still “statistically significant?” These are the broader issues you will address here.

Rosenberg, Reid, and Lanstein and Fama and French distinguish value from growth stocks by comparing each company's "book value" to its "market value." Book value measures the accounting value of the firm's assets per share. Market value is just the market price of each share of stock. According to Rosenberg and Fama and French, therefore, value stocks are those with high book compared to market values, and growth stocks are those with low book to market values.

Let \tilde{r}_H be the random return on a portfolio of high book-to-market value stocks, and let \tilde{r}_L be the random return on a portfolio of low book-to-market growth stocks. As usual, let \tilde{r}_M be the random return on the market portfolio and let r_f be the risk-free rate. Then, as we know, the CAPM implies

$$E(\tilde{r}_H) = r_f + \beta_H[E(\tilde{r}_M) - r_f] \quad (1)$$

and

$$E(\tilde{r}_L) = r_f + \beta_L[E(\tilde{r}_M) - r_f], \quad (2)$$

where β_H and β_L are the CAPM betas on the portfolios of value and growth stocks.

To concentrate specifically on the returns from a value investing strategy, Fama and French define an HML ("high-minus-low") portfolio that takes long positions in value stocks and short positions of equal value in growth stocks. The random return on this HML portfolio is therefore

$$\tilde{r}_{HML} = \tilde{r}_H - \tilde{r}_L.$$

and, according to the CAPM

$$E(\tilde{r}_{HML}) = \beta_{HML}[E(\tilde{r}_M) - r_f], \quad (3)$$

where

$$\beta_{HML} = \beta_H - \beta_L$$

is the CAPM beta on the HML portfolio, and also the difference between the beta on the value portfolio and the beta on the growth portfolio.

What is striking is that the risk-free rate r_f gets eliminated when (3) is derived by subtracting (2) from (1). This happens because the HML portfolio attempts to generate an "excess return" by taking long and short positions of equal value in the two sets of stocks: value (long) and growth (short). This implication of (3) will form the basis of our statistical test of the CAPM and, specifically, of the continued profitability of the value investing strategy embodied in Fama and French's HML portfolio.

Consider, in particular, a regression of the return on the HML portfolio on a constant and the difference between the market return and the risk-free rate:

$$\tilde{r}_{HML,t} = \alpha + \beta(\tilde{r}_{M,t} - r_{f,t}) + e_t \quad (4)$$

where $\tilde{r}_{HML,t}$ is the HML return at time t , $\tilde{r}_{M,t}$ and $r_{f,t}$ are the market return and the risk-free rate at time t , and e_t is the regression error term. Remember from econometrics that the regression error term has mean (expected value) zero: $E(e_t) = 0$ Therefore, (4) implies

$$E(\tilde{r}_{HML,t}) = \alpha + \beta[E(\tilde{r}_{M,t}) - r_{f,t}]. \quad (5)$$

Comparing (3) and (5), we can see that the slope coefficient β from the regression equation (4) provides an estimate of the HML portfolio's CAPM beta. More importantly for our purposes, (3) implies that according to the CAPM, the intercept term α in the regression equation (4) should equal zero. We can test the CAPM's implications, therefore, by seeing whether the estimated value of α in (4) is statistically different from zero. The annual data set mentioned above contains data on the HML return $\tilde{r}_{HML,t}$ and the excess return on the market $\tilde{r}_{M,t} - r_{f,t}$ that you can use to run this statistical test. Since, for this question, we are interested in interpreting the past as opposed to predicting the future, it makes sense to use the full range of annual data running from 1927 through 2020.

Use these data to estimate the regression in (4). Then, use the t -statistic to test the null hypothesis, implied by the CAPM, that the intercept is equal to zero.

There are at least two ways of doing this. Choose which is easiest or most convenient for you.

One way is to load the data into a statistics or econometrics software package, like Stata or R, and regress the HML return on a constant and the excess return on the market. Most software packages will report the estimated intercept and slope coefficients, their standard errors, and the t -statistic used to test the null hypothesis that the intercept (or the slope, but here we're interested in the intercept), equals zero. Remember that for data samples of the size we have here, with $n = 94$ annual observations covering 1927-2020, a t -statistic that is larger than 1.99 means that you can reject the null hypothesis at the 95 percent confidence level and a t -statistic that is larger than 1.66 means that you can reject the null hypothesis at the 90 percent confidence level.

A second way is to use Google Sheets. This requires several steps, but none of these steps is terribly onerous; so you can do this if you don't want to re-familiarize yourself with whatever statistics software package you may have used for stats and/or econometrics.

The first step is to download the Excel spreadsheet and open it, instead, with Google Sheets. Then, estimate the intercept from the regression in (4) using the command

```
=intercept(e11:e104,d11:d104)
```

and the slope using the command

```
=slope(e11:e104,d11:d104)
```

You should find that the slope coefficient is positive but relatively small. What this tells us is that value stocks do have CAPM betas that are slightly higher than those of growth stocks. Some of the higher returns provided historically by value stocks over growth stocks, therefore, can be explained by the CAPM as compensation to investors for taking on more aggregate risk.

The difference between CAPM betas, however, is much too small to explain all of the difference. You should also find that the estimated intercept term is slightly less than 3.5.

What this tells us is that, even after accounting for differences in CAPM betas, value stocks have, over the past 94 years, generated average annual returns that are nearly 3.5 percentage points higher than growth stocks.

The final question is whether the estimate of the intercept term is large enough, relative to its standard error, to be able to say that it is “statistically significant.” Unfortunately, there’s no simple command (at least that I know of) that computes standard errors in Google Sheets. So you will need to do this manually.

Start by creating a new column in the Google sheet that will contain the fitted values of the regression error term e_t from (4). Then, fill in the values using the formula

$$\hat{e}_t = \tilde{r}_{HML,t} - \hat{\alpha} - \hat{\beta}(\tilde{r}_{M,t} - r_{f,t}),$$

where $\tilde{r}_{HML,t}$ and $\tilde{r}_{M,t} - r_{f,t}$ are, again, the series in columns E and D and where $\hat{\alpha}$ and $\hat{\beta}$ are your intercept and slope estimates from before. Then (assuming you put this series for \hat{e}_t in column F of the Google sheet) you can compute the variance of the regression errors using the command

`=var(f11:f104)`

Likewise, compute the mean and variance of the independent variable, the excess return on the market, using the commands

`=average(d11:d104)`

and

`=var(d11:d104)`

Finally, let σ_e^2 denote the variance of the regression errors, let \bar{x} be the mean of the excess return on the market, and let σ_x^2 be the variance of the excess return on the market; those are the three statistics you just finished computing in the previous steps. If you search through enough econometrics textbooks, you’ll eventually find that the standard error for $\hat{\alpha}$ equals

$$SE(\hat{\alpha}) = \sqrt{\left(\frac{\sigma_e^2}{n}\right) \left[\frac{\sigma_x^2 + (\bar{x})^2}{\sigma_x^2}\right]},$$

where $n = 94$ is the sample size. Note that to apply the formula, you’ll need to square the mean \bar{x} of the market return when computing the second fraction and also take the square root of the product of the two fractions. Now recall that the t -statistic

$$\frac{\hat{\alpha}}{SE(\hat{\alpha})}$$

obtained by dividing $\hat{\alpha}$ by its standard error rejects the null hypothesis that $\hat{\alpha} = 0$ if it exceeds 1.99 for 95 percent confidence or 1.66 for 90 percent confidence.

So can we still say with statistical confidence that a value investing strategy provides expected returns above those predicted by the CAPM, even after the past 14 years when growth stocks have performed much better? Once you’ve answered this question you will know!