Final Exam

ECON 337901 - Financial Economics — Peter Ireland
Boston College, Department of Economics — Spring 2026

Due Monday, May 11, 12noon

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 2021 through 2025, is available either as an Excel spreadsheet or a plain text file at

`http://irelandp.com/econ3379/exams/final26s_monthly.xlsx`

or

`http://irelandp.com/econ3379/exams/final26s_monthly.txt`

The other, containing annual data on stock returns from 1927 through 2025, is available in the same two formats at

`http://irelandp.com/econ3379/exams/final26s_annual.xlsx`

or

`http://irelandp.com/econ3379/exams/final26s_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 2021 through December 2025 on the market as a whole and on six individual stocks: Advanced Micro Devices (a manufacturer of computer chips), Kraft Heinz (food and beverages), Netflix (streaming video), Nucor (steel and steel products), AT&T (telecommunications), and Exxon-Mobil (oil and natural gas). **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 Advanced Micro Devices' 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 stock returns, in particular, have been 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 2025. Because this period both includes the Great Depression 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 years from 2000 through 2025. Over that period, the average return on the stock market has been just slightly less than 10 percent, which suggests setting $E(\tilde{r}_M) = 10$ in the CAPM formula.

Meanwhile, although the Federal Reserve is presently targeting interest rates in a range between 3.50 and 3.75 percent, most financial market participants expect that the Fed will reduce interest rates at least somewhat later this year. Based on these considerations, let's take $r_f = 3$ as our best estimate for the risk-free rate in the CAPM formula.

Thus, plugging $E(\tilde{r}_M) = 10$ and $r_f = 3$ together with your estimates of the betas from question 1 into the right-hand side of the CAPM formula

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

will produce the CAPM estimates of the expected return on each of the six stocks.

3. Momentum 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, using a strategy known as “momentum investing.” As its name suggests, this strategy involves buying stocks (“winners”) that have gone up over the past year and avoiding stocks (”losers”) that have gone down over the past year.

The superior performance this strategy can be quantified using the “winners-minus-losers” or “WML” measure developed by Eugene Fama and Kenneth French and described in more detail below. According to this measure, the previous year’s winners provided annual returns that are more than 8.5 percentage points higher than the returns on the previous year’s losers over a long period extending from 1927 through 2025.

The difference between returns on winners versus losers is very volatile, however. In 1980, for example, returns on winners exceeded those on losers by more than 36 percentage points. In 2009 – the worst year ever for this investment strategy – returns on losers exceeded those on winners by more than 83 percentage points! In fact, the strategy delivered returns that are less extreme but still quite disappointing in 2003, 2016, and 2023; in each of those years, returns on losers exceeded those on winners by more than 20 percentage points.

Of course, financial economists will always be quick to point out that higher average or expected returns on individual stocks or portfolios of individual stocks do not necessarily violate the implications of the CAPM. Suppose, in particular, that winners have higher CAPM betas than losers. Then the CAPM predicts that winners will have higher expected returns as well. The interpretation of the momentum investment strategy would then be: it provides higher expected returns, but only because it exposes the investor to more aggregate risk.

It turns out, though, that two famous research articles, one by Narasimhan Jegadeesh and Sheridan Titman (“Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *Journal of Finance*, March 1993) and the other by Eugene Fama and Kenneth French (“Size, Value, and Momentum in International Stock Returns,” *Journal of Financial Economics*, September 2012), find that higher average returns on winners versus losers *do* violate the CAPM. That is, winners offer higher average returns than losers, even after accounting for differences in their CAPM beta.

By answering this question, you will see whether or not these earlier results continue to hold in the longer set of annual data, running from 1927 through 2025, mentioned above. This data set runs well beyond those used in the original studies and even includes the terrible year 2009 for momentum investing, 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.

Let $\tilde{r}_W$ be the random return on a portfolio of winners (again, stocks that have gone up a lot in the past year) and let $\tilde{r}_L$ be the random return on a portfolio of losers (stocks that have gone down a lot in the past year). 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

$$E(\tilde{r}_W) = r_f + \beta_W[E(\tilde{r}_M) - r_f] \tag{1}$$

and

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

where β_W and β_L are the CAPM betas on the portfolios of winners and losers.

To concentrate specifically on the returns from following the momentum strategy, Fama and French define their WML portfolio as one that takes a long position in the portfolio of winners and a short position of equal value in the portfolio of losers. The random return on this WML portfolio is therefore

$$\tilde{r}_{WML} = \tilde{r}_W - \tilde{r}_L.$$

and, according to the CAPM,

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

where

$$\beta_{WML} = \beta_W - \beta_L$$

is the CAPM beta on the WML portfolio, and also the difference between the beta on the portfolio of winners and the beta on portfolio of losers.

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 WML portfolio attempts to generate an "excess return" by taking long and short positions of equal value in the two portfolios: of winners (long) and losers (short). This implication of (3) will form the basis of our statistical test of the CAPM and, specifically, of the continued profitability of the momentum trading strategy.

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

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

where $\tilde{r}_{WML,t}$ is the WML 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

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

Comparing (3) and (5), we can see that the slope coefficient β from the regression equation (4) provides an estimate of the WML 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 WML return $\tilde{r}_{WML,t}$ 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 2025, which includes some very good years and very bad years for the momentum strategy.

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 WML 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 $n = 99$ annual observations covering 1927-2025, 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:e109,d11:d109)
```

and the slope using the command

```
=slope(e11:e109,d11:d109)
```

You should find that the slope coefficient is actually negative, though small. What this tells us is that the success of the momentum strategy cannot be explained by differences in CAPM betas across winners and losers: according to the CAPM, *losing* stocks should have slightly higher expected returns!

More important, you should also find that the estimated intercept term is positive and large: equal to approximately 9.4. What this tells us is that, after accounting for differences in CAPM betas, the portfolio of winning stocks has provided average annual returns that are more than 9.4 percentage points greater than the returns on the portfolio of losing 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

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

where $\tilde{r}_{WML,t}$ 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:f109)
```

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

```
=average(d11:d109)
```

and

```
=var(d11:d109)
```

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

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

where $n = 99$ 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

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

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 the momentum strategy provides expected returns above those predicted by the CAPM? Once you've answered this question you will know!