

Problem Set 1

Due in class, week 1

Do the readings, as specified in the syllabus. Answer the following problems.

Note: in this and following problem sets, make sure to answer the direct questions, as well as make the indicated plots and tables. That's how we grade them.

Part I

These problems should help you to get up to speed with time series. Read the time series section of the notes first if this isn't easy.

1. A stock price x_t follows an AR(1) process, i.e.

$$x_t = \rho x_{t-1} + \varepsilon_t$$

where ε_t are a sequence of i.i.d. normally distributed random variables with mean 0 and variance σ_ε^2 .

- (a) Find the sequence of conditional means. I.e., if you have observed values $\{x_t, x_{t-1}, \dots\}$ find $E_t(x_{t+1}), E_t(x_{t+2}), E_t(x_{t+j})$ in terms of those observed values. (Hints: You may have seen the conditional expectation $E_t(x_{t+j})$ written as $E(x_{t+j}|x_t, x_{t-1}, \dots)$ or $E(x_{t+j}|I_t)$ I.i.d. means that $E_t(\varepsilon_{t+j}) = 0, \sigma_t(\varepsilon_{t+j}) = \sigma_\varepsilon^2$. Start by writing $x_{t+1} = \rho x_t + \varepsilon_{t+1}; x_{t+2} = \rho^2 x_t + \rho \varepsilon_{t+1} + \varepsilon_{t+2}, \dots$ and then take means and variances given information available at time $t, \{x_t, x_{t-1}, \dots\}$.)
- (b) Find the sequence of conditional variances. $\sigma_t^2(x_{t+1}), \sigma_t^2(x_{t+2})$ etc. (Again, you may have seen conditional variance written as $\sigma^2(x_{t+j}|x_t, x_{t-1}, \dots)$ or $\sigma^2(x_{t+j}|I_t)$.)
- (c) Suppose $\rho = 0.95, \sigma(\varepsilon) = 0.25, x_t = 0$ for all $t \leq 0$, and at time 1 you see $x_1 = 1$.
 - i. Plot the conditional mean, where we expect x_t to go in the future given we have seen $x_1 = 1$, i.e. plot $E_1(x_t)$ vs. t for $t = 1, 2, \dots, 20$.
 - ii. Add plus and minus one conditional standard deviation bands, i.e. plot $E_1(x_t) \pm \sigma_1(x_t)$ for $t = 1, 2, \dots, 20$. This plot gives you a sense of how uncertain the forecast is, the range of things that might happen. Risk management is as (if not more) important than forecasting!
 - iii. Add a simulation of x_t to your plot – pick $\{\varepsilon_t, t = 2, 3, 4, \dots\}$ from a random number generator and plot the resulting path of x_t . ($\varepsilon_1 = 1$ of course.) (In matlab, the `randn` command makes random normals. Then you can use `for i=1:20, ... end` to loop through and make x from ε). You should do this several times for yourself and watch how different random numbers lead to different paths.
The point of this is to get a feeling for how x_t “mean-reverts” and how uncertainty about its future builds up through time. By looking at the simulated values you should get some sense of how the conditional mean (forecast) and variance relate to the actual series.
- (d) Find the *unconditional* mean and variance $E(x_t)$ and $\sigma^2(x_t)$. There are two ways to do this. First, the unconditional mean is the limit of the conditional mean as you move information backwards through time, $E(x_t) = \lim_{j \rightarrow \infty} E_{t-j}(x_t)$. Second (better), you can just take the unconditional expectation of both sides of $x_t = \rho x_{t-1} + \varepsilon_t$ yielding

$$E(x_t) = \rho E(x_{t-1}).$$

Then, realize that the mean of x_t and x_{t-1} are the same – unconditional moments don't depend on calendar time.

(e) How would you generalize the AR(1) model $x_t = \rho x_{t-1} + \varepsilon_t$ to produce a mean $E(x_t) = \mu$?

2. In this problem, I want you to think about how returns look at different time horizons, and how predictability of returns might change that. Suppose first that one-year log returns $r_t = \log(R_t)$ are not correlated over time $cov(r_t, r_{t+j}) = 0$ and have mean $E(r_t) = \mu$ and variance $\sigma^2(r_t) = \sigma^2$ that are constant over time. The compound log long-horizon return is $\log(R_t R_{t+1} R_{t+2} \dots R_{t+k}) = r_{t+1} + r_{t+2} + \dots + r_{t+k}$ and the *annualized* compound log long-horizon return is $\log \left[(R_t R_{t+1} R_{t+2} \dots R_{t+k})^{\frac{1}{k}} \right] = \frac{1}{k} [r_{t+1} + r_{t+2} + \dots + r_{t+k}]$. Now, let's think about how these returns scale with horizon. The big question underlying this analysis is, are returns in some sense "safer" for long-horizon investors?

(a) Find the mean log long-horizon return $E[r_{t+1} + r_{t+2} + \dots + r_{t+k}]$ and the mean annualized log long horizon $E\left(\frac{1}{k} [r_{t+1} + r_{t+2} + \dots + r_{t+k}]\right)$ return as a function of the horizon k and the one-period mean μ . If the mean annual return is 6%, what is the mean monthly return? What is the mean annualized monthly return?

(b) Find the variance of the log long-horizon return $\sigma^2[r_{t+1} + r_{t+2} + \dots + r_{t+k}]$ and of the annualized log long horizon return $\sigma^2\left(\frac{1}{k} [r_{t+1} + r_{t+2} + \dots + r_{t+k}]\right)$ as a function of the horizon k and the one-period variance σ^2 .

(c) Find the standard deviation of the log long-horizon return $\sigma[r_{t+1} + r_{t+2} + \dots + r_{t+k}]$ and of the annualized log long horizon return $\sigma\left(\frac{1}{k} [r_{t+1} + r_{t+2} + \dots + r_{t+k}]\right)$ as a function of k . If the standard deviation of annual returns is 16%, what is the standard deviation of monthly and daily returns?

(d) A stock has mean and standard deviation of monthly returns equal to 1%. Express these on an annualized basis.

(e) Find the "Sharpe ratio" (quotes since you're using logs and not subtracting a risk free rate) $E(r_{t+1} + \dots + r_{t+k})/\sigma(r_{t+1} + \dots + r_{t+k})$ as a function of horizon k . Does the Sharpe ratio depend on the units (annual vs. monthly return) in which you quote it?

(f) The typical portfolio allocation formula says stock allocation should depend on the ratio of mean to *variance* of total return:

$$\text{stock share} = \frac{1}{\text{risk aversion}} \times \frac{E(r)}{\sigma^2(r)}.$$

For example, if mean returns are 8% (0.08), the standard deviation of returns is 20% so $\sigma^2(r) = 0.04$, and risk aversion is 2, then the investor should put all his money in stocks, share = 1. Now, how does this advice scale with horizon? Should an investor with a 10 year horizon allocate more to stocks than an investor with a 1 year horizon because "stocks are safer in the long run" and he can "wait out market declines"?

(g) Client: "No, you have it wrong. The mean return is independent of horizon, but standard deviation of returns is lower at long horizons because of time-diversification. Long horizon average returns are very stable, and volatility of long-run average returns goes down with the square root of horizon. I really can afford to take more equity risk because I'm a long horizon investor." In the context of the above formula, (and our assumption that returns are independent over time) explain the mistake he is making.

3. Now, what if returns have some "momentum" or "mean-reversion" in stock returns? Does this make "long horizon returns safer?" To answer this question, suppose returns follow the MA(1) process $r_t = \mu + \varepsilon_t + \theta \varepsilon_{t-1}$ with ε_t independent over time and with $E(\varepsilon_t) = 0$ and $\sigma^2(\varepsilon_t) = \sigma_\varepsilon^2$

- (a) Start by getting familiar with the MA(1).
- Find $E(r_t)$, $\sigma^2(r_t)$, and $cov(r_t, r_{t\pm 1}), cov(r_t, r_{t\pm 2})$. (“Find” means this stuff on the left and $\mu, \theta, \sigma_\varepsilon^2$ on the right.) How is the pattern of covariances differ from the AR(1) case?
 - Plot and interpret the impulse-response function of r_t , i.e. how it behaves if hit by $\varepsilon_t = 1$ and all other $\varepsilon = 0$. Also find and plot the impulse-response function of the return and “price” or *cumulative* log return $p_t = r_0 + r_1 + r_2 + \dots + r_t$.
In other words, take $r_{-1} = r_0 = \varepsilon_{-1} = \varepsilon_0 = 0$. Then suppose $\varepsilon_1 = r_1 = 1$, and all future $\varepsilon = 0$. Plot where r_t goes in response to this shock. Make a plot for $\theta = +0.5$ (“momentum”) $\theta = 0$ (the iid return benchmark) and $\theta = -0.5$ (“mean=reversion”).
Point: This plot corresponds to the expected value plot you made for the AR(1) above. It should give you a feeling for the “persistence” or “mean-reversion” nature of the MA(1) and how it captures different behavior from the $\theta = 0$ benchmark case from problem 2.
- (b) Now find the mean $E[r_{t+1} + r_{t+2} + \dots + r_{t+k}]$ and variance $\sigma^2[r_{t+1} + r_{t+2} + \dots + r_{t+k}]$ of long horizon returns in terms of the mean $E(r) = \mu$ and variance $\sigma^2(r)$ of one year returns. Does mean-reversion make long horizon returns more/less attractive than one year returns? How about momentum? (Note returns are *not* uncorrelated over time here. $cov(r_t, r_{t+1}) \neq 0$. There is more to this than just restating problem 2!)

Point: This is really the big payoff moment, so stop to appreciate it. You discover in problem 2 that when returns are independent over time, stocks are not safer for long run investors. This is a classic theorem in finance and contrary to usual investment advice. How could this be wrong? Well, if there is a tendency for stocks to “bounce back” after declines, then they really are safer for long run investors. The $\theta < 0$ model here has some “bounce back” in prices, as you see from the impulse response graph, and you verify in the last question that if it is right, stocks really are “safer for long run investors.” Once again, when our models don’t fit the standard advice, sometimes it’s the model that’s wrong, not Wall Street!

Part II

This set of problems is a regression review, and it’s also important background for the first week discussion. (We’re going to run a lot of regressions this quarter!) Load the data from the class website. This data comes from CRSP via WRDS. The stock return is the return on the value-weighted portfolio of all US stocks – NYSE+AMEX+NASDAQ. The treasury bill return is the return for the year on 90 day treasury bills.

- Make a plot of the..
 - dividend/price and price/dividend ratios,
 - returns and treasury bill returns over time, in annual percents.

Make sure these plots look reasonable, and that you understand the units of the variables—are the numbers about right?

- Run regressions of returns at date $t + 1$ on returns at date t , i.e.

$$R_{t+1} = a + bR_t + \varepsilon_{t+1}; t = 1, 2, \dots, T - 1.$$

Do this for

- The stock return
- The treasury bill return

(c) The stock excess return, i.e. stock return - treasury bill return.

Make a table of the coefficients, t statistics, and R^2 for these regressions. Include the mean return $E(R)$ and the standard deviation of expected returns $\sigma(E_t(R_{t+1})) = \sigma(a + bR_t)$. Express everything appropriate in annual percent units¹.

Regression of returns on lagged returns

$$R_{t+1} = a + bR_t + \varepsilon_{t+1}$$

	b	t(b)	R^2	E(R)	$\sigma(E_t(R_{t+1}))$
Stock	0.01			11.62	0.13
T bill	0.92				
Excess					

Note: The standard deviation of *expected* returns is a subtle concept. It's the numerator of $R^2 = \sigma^2(a + bR_t)/\sigma^2(R_t)$. Here, it's interesting to compare $\sigma(a + bR_t)$ to the *mean* of R_{t+1} , not to the *variance* of R_{t+1} . This answers the question, "how much to mean returns change over time?" as opposed to "how much of the variation in returns is predictable ahead of time?" The answer to the former can be "very big" even if the latter is "small."

Hint: to run this regression, your left hand variable is the return from 1927 to the end of the sample. Your right hand variables are a constant (column of ones) and the return from 1926 to one less than the end.

$$y = \begin{bmatrix} R_{1927} \\ R_{1928} \\ \dots \\ R_{2009} \end{bmatrix}; x = \begin{bmatrix} 1 & R_{1926} \\ 1 & R_{1927} \\ 1 & \dots \\ 1 & R_{2010} \end{bmatrix};$$

The single greatest problem people have with this regression is lining up the data right. (Whether you include the column of ones depends on your regression program. Many programs add the column for you. If you use my regress_jc.m, they you need the extra column.)

- Run regressions of excess returns at date $t + 1$ on the dividend-price ratio at date t , i.e.

$$R_{t+1}^{\text{stock}} - R_{t+1}^{\text{treas.bill}} = a + b(D/P)_t + \varepsilon_{t+1}; t = 1, 2, \dots, T - 1.$$

Compare the results to the excess return regression from part 2c. You should see a dramatic difference. *This is the big point of the first class.*

- Run regressions of 5 year excess returns on the dividend price ratio. To do this, first form 5 year compound stock and treasury bill returns,

$$R_{t,t+5}^{\text{stock}} = R_{t+1}^{\text{stock}} \times R_{t+2}^{\text{stock}} \times R_{t+3}^{\text{stock}} \times R_{t+4}^{\text{stock}} \times R_{t+5}^{\text{stock}}$$

and the same for bonds (get the units right - R must be a number like 1.10 for you to be able to compound it, not a number like 0.10 or 10). Then the excess returns is

$$R_{t,t+5}^{\text{Excess}} = R_{t,t+5}^{\text{stock}} - R_{t,t+5}^{\text{bond}}$$

Run

$$R_{t,t+5}^{\text{Excess}} = a + b(D/P)_t + \varepsilon_{t+5}$$

¹ "Standard deviation of *expected* returns is a subtle concept. It's the numerator of $R^2 = \sigma^2(a + bR_t)/\sigma^2(R_t)$. Here, it's interesting to compare $\sigma(a + bR_t)$ to the *mean* of R_{t+1} , not to the *variance* of R_{t+1} . This answers the question, "how much to mean returns change over time?" as opposed to "how much of the variation in returns is predictable ahead of time?" The answer to the former can be "very big" even if the latter is "small."

- (a) Use overlapping observations. Your first data point is the 1926-1931 return (i.e. $R_{1927} \times R_{1928} \times R_{1929} \times R_{1930} \times R_{1931}$) on the 1926 D/P; your second data point is the 1927-1932 return on the 1927 D/P, etc.
- (b) The overlap is a bit weird. Does this regression correspond to the OLS assumptions? Is the regression coefficient biased? Are the standard errors ok? Is the R^2 biased? (Hint: Ok, what were the OLS assumptions, and which assumptions did you need for unbiased coefficient and which did you need for standard errors?)
- (c) Use nonoverlapping observations. Your first data point is the 1926 to 1931 return ($R_{1927} \times R_{1928} \times \dots \times R_{1931}$) on the 1926 D/P; your second data point is the 1931 to 1936 ($R_{1932} \times R_{1933} \times \dots \times R_{1936}$) return on the 1931 D/P, etc. Compare the results to the part a results. Did the coefficients change a lot? Did the standard errors and t statistics change a lot? (This is a simple technique for getting standard errors right. We'll talk about fancier ones later.)
- (d) Make a plot of the actual 5 year return and the prediction, i.e. fitted value of the regression from part a. Align the data so that the prediction and return are at the same time, e.g. plot the 1926-1931 return at 1931, and plot $a + b(D/P)_{1926}$ also at 1931.
- (e) Compare the 5 year regression results with the one-year regression results from question 3. How does the prediction of 5 year returns differ from one-year returns?

Please do not turn in reams of unformatted regression output! Make little tables of the numbers that matter.

If you're using matlab and new to it, check out the matlab hints for problem set 1 below. Also I have written a small enhancement to the matlab function regress.m, called regress_jc.m, on the class website. Feel free to use it to run the regressions.

The number one problem with getting things like this to work is confusing the units of returns and growth rates. Make very sure you know whether you're using net returns r , numbers like 0.1; when you're using gross returns R , numbers like 1.1, when you're using percentages, numbers like 10, and when you're using log or continuously compounded returns, also numbers like 0.1.

Matlab hints for Problem Set 1

The help screen should pop up with “begin here”. If it doesn’t, type `>>helpdesk [enter]` and it will pop up. Read “begin here” then “getting started.” Read “introduction” “matrices and arrays” “graphics” and “programming.” You do not need “creating graphical user interfaces” or “desktop tools.” (You can also read this stuff in the printed version, but if you do it on the computer you can try it as you go, which is much more effective. The printed version is available from the help system under “printable documentation.”)

Mathworks has a nice tutorial here: www.mathworks.com/academia/matlabtutorial

General:

In the main window, use the next to "current directory" to get yourself to a comfortable place to work. Put the data file in this directory. Do the problem in an m file, not at the command line. To do that,

```
>>edit ps1.m
```

will get you started. Write, save the file, and then

```
>> ps1.m
```

will run it.

I start all my m files with

```
clear all;  
close all;
```

This closes up graphics windows and makes sure I am starting with a fresh workspace.

There are two help systems. If you need quick help on a specific command (syntax etc.) you can get it quickly with

```
>> help commandname
```

If you need to find a command to do something or want more serious documentation (useful for plot, for example) bring up the help system with

```
>> helpdesk
```

(if it’s not already up) and use the search tool or the index.

Some installations of the student edition get stuck looking for the license file. I hear the following tip works.

Add the following into the property of the shortcut to launch matlab

```
"C:\MATLAB_SV701\bin\win32\MATLAB.exe -c C:\MATLAB_SV701\bin\win32\license.dat"
```

The key is to add "-c C:\MATLAB_SV701\bin\win32\license.dat" when launching matlab. Path needs to be changed according person’s own installation path.

Some useful commands for problem set 1: (use help for more details)

`x = load('filename.txt');` loads text files (ps1.txt); assigns the results to matrix x.

`x = filename(:,2);` sets variable x equal to the second column of numbers in filename

`a.*b; a./b;` does an element by element multiplication and division of matrix a and matrix b

`T = size(x,1);` creates variable T; if x has 3 rows and 2 columns, T=3.

`rhv = [ones(T,1) x]` creates matrix with one column of ones and then x;

`x1 = x(1:T-1,:)` creates matrix missing the last observation of x, i.e. 1926 to 2002

`y = x(2:T,:);` creates matrix missing the first observation, i.e. from 1927 to 2003

`[b,bse,bt,r2,F] = regress_jc(lhv,rhv,display);` invokes the program regress_jc (must reside in same directory). Returns b,bse,bt,r2,F to current directory.

`b = x\y` run regression of y on x. (To include a constant, the first column of x should be all ones). Use this if you want to write your own regression programs

Tables

disp and format are easy ways to show numbers, e.g.

`format bank` prints all numbers in two decimal format

`disp('Hello Kitty');` prints header for table

`disp(x);` prints matrix x

`fprintf` makes really pretty tables but is much harder to use.

Plots:

`(1926:2013);` creates a row vector (1926 19272002 2013). Useful for plotting

`figure;` opens new figure window

`plot((1926:2013)',x);` plots time series in column vector x (see plot documentation for lots of options)

`title('Dividend price ratio');` title of plot

`axis([1926 2013 -inf inf]);` scales axes; here from 1926 to 2013, and automatically choose vertical scale

`legend('Stock','T Bill',h);` puts legend on plot, in position h (can omit h; see help legend)

`print -format filename;` automatically "prints" figure to file so you can import it to your program