

Problem Set 2

Due by Sat 12:00, week 2

Part I. Reading questions: These refer to the reading assignment in the syllabus. Please hand in *short* answers. Where appropriate, cite the text or tables for your answers

Short (*short!*) answer questions to review this week's reading.

1. If the dividend yield of the market portfolio rises from 3% to 4%, how much (roughly) does your forecast of the market return over the next year change (in percentage points)?
2. Suppose $R_{t+1} = a + bx_t + \varepsilon_{t+1}$; $a = 0$; $b = 1$; $x_t = 1$; $\sigma(x) = 0.3$; $\sigma(\varepsilon) = 0.4$. What word describes each quantity, and what is the value of each one?
 - (a) $E(R_{t+1})$
 - (b) $E_t(R_{t+1})$
 - (c) $\sigma_t(R_{t+1})$
 - (d) $\sigma[E_t(R_{t+1})]$
 - (e) $\sigma(R_{t+1})$
 - (f) R^2 of this regression

(Note: I picked numbers that would be easy for the calculation, not realistic.)

3. If expected returns rise, which of the following should happen (and what's wrong with the other idea?)
 - (a) Traders seeing the higher expected returns will flock to buy stocks, sending the prices up.
 - (b) Higher expected returns reflect general fear. People trying to sell drives prices down.
4. For a one-period asset, that satisfies $R_{t+1} = \frac{D_{t+1}}{P_t}$, suppose we run a regression of log returns and log dividend growth on the dividend yield,

$$\begin{aligned} r_{t+1} &= a_r + b_r(d_t - p_t) + \varepsilon_{t+1}^r \\ \Delta d_{t+1} &= a_d + b_d(d_t - p_t) + \varepsilon_{t+1}^d \end{aligned}$$

what would we conclude if we were to find that $b_r = 0$, returns are unpredictable, $b_d = 0$, dividend growth is unpredictable, and $\sigma(d - p) = 0.15$ (roughly its value in data)?

5. Consider the identity

$$p_t - d_t \approx \sum_{j=1}^{\infty} \rho^{j-1} (\Delta d_{t+j} - r_{t+j}).$$

How can this identity have something known at time t on the left, but things known only after time t on the right? (Hint: If you're confused think about the one period version first.)

Part II Computer

We'll start by replicating the VAR (vector autoregression) and impulse-response function. We'll use the full 1926-now CRSP database so that the answers don't look exactly as they do in the class notes, and (minor) to see if that data makes a big difference in results. Use `ps1_data.txt` again.

1. The identity

$$r_{t+1} \approx -\rho dp_{t+1} + dp_t + \Delta d_{t+1} \quad (1)$$

is only approximate. The rest of the problem will work better if we use data that really do conform to the identity. To this end, compute a dividend growth series that satisfies the identity exactly,

$$\Delta d_{t+1} = \rho dp_{t+1} - dp_t + r_{t+1}.$$

Here and throughout the problem set, use $\rho = 0.96$. Make a plot comparing this to the actual log dividend growth series. You should see that the approximation is excellent, so you really don't lose much at all by using this new dividend growth series rather than the original one.

(There is a constant in (1) which we will ignore, as it only affects constants in regressions. Thus you can ignore any difference in level between series, or plot and use series after taking out their mean. If you want to dig deeper, use the version of (1) with the constant.)

2.

(a) Run forecasting regressions of returns, this new dividend growth and dividend yield on the dividend yield

$$\begin{aligned} r_{t+1} &= a_r + b_r dp_t + \varepsilon_{t+1}^r \\ \Delta d_{t+1} &= a_d + b_d dp_t + \varepsilon_{t+1}^d \\ dp_{t+1} &= a_{dp} + b_{dp} dp_t + \varepsilon_{t+1}^{dp} \end{aligned}$$

(b) What does the identity (1) imply for the relationship between regression coefficients b_r, b_d, b_{dp} ? (Hint: substitute the right hand side of the regression equations in to (1)) Calculate what b_r should be given b_d and b_{dp} and compare to your actual estimated b_r . How close is the prediction?

(c) What does the identity (1) imply for the relationship between the residuals $\varepsilon_{t+1}^r, \varepsilon_{t+1}^d, \varepsilon_{t+1}^{dp}$? Check to see how well this relationship holds in your estimates. (See how close the ε_{t+1}^r you calculate from ε_{t+1}^d and ε_{t+1}^{dp} is to the actual ε_{t+1}^r .) Report your first 5 observations of the actual ε_{t+1}^r and the value imputed from $\varepsilon_{t+1}^d, \varepsilon_{t+1}^{dp}$ and the identity.

(d) Repeat parts a and b using the real, original, dividend growth series. How different is b_d ? How close is the relationship between the b coefficients? For the rest of the problem use the implied dividend growth series so that you can see the identities at work.

(e) Sum up in a few words: What is the implication of the return identity (1) for the three regressions in 2a? Why, as intuitively as you can?

3. Using your regression estimates, plot responses of returns, cumulative returns, dividend growth, cumulative dividend growth, and dividend yield to

(a) the dividend yield shocks $\varepsilon^{dp} = 1, \varepsilon^d = 0$ and

(b) the dividend growth shocks, $\varepsilon^{dp} = 0, \varepsilon^d = 1$.

Make sure there is a contemporaneous ε^r shock in both cases so that the identity (1) is satisfied. (I also use the identity

$$\Delta p_{t+1} = -dp_{t+1} + dp_t + \Delta d_{t+1}$$

to find the path of prices, and I plotted the level of dividends and prices too. These are welcome, but optional.) Unlike the simple version in class or notes, use your estimated $b_d \neq 0$ coefficient.

- (c) Why does the return response in the period of the shock look so different from the return response in the periods after the shock?
4. What would a finance professor from the 1970s expect for this regression and impulse-response function? He would certainly expect $b_r = 0$ – returns are not forecastable. He would not be surprised at your sample value for b_{dp} , as he knew dividend yields move slowly over time. But...
- (a) What would he expect for b_d and why? Answer in equations in words and in numbers, using $b_r = 0$ and your estimate for b_{dp} . (Hint: is it possible that both returns and dividend growth are unforecastable? How would he answer an efficiency critic who said “prices are way high relative to dividends, there must be a bubble going on?”)
- (b) Use $b_r = 0$, $b_{dp} =$ unchanged from your estimate, and your value for b_d from part a to plot new impulse-response functions to a dp shock and a Δd shock. These are “What the finance professor from the 1970s expected to see.” (Mechanically, just reuse the code from the last question with these new coefficients. It takes longer to talk about than to do it.)
- (c) Explain the parts of the impulse-response graphs that look the same and the parts of the graph that look different
- (d) In our original VAR, I called the two shocks “dividend growth shock” (with no change in expected returns) and “expected return shock” (with no change in dividend growth), based on the pattern of impulse-responses. What are good names for the shocks in your new pictures? (The point of this question is to think about ir functions “backwards,” i.e. “what news about future dividends and returns made prices change as they did (or did not)?”
5. Suppose expected returns $x_t = E_t(r_{t+1})$ follow an AR(1). This means

$$\begin{aligned}x_t &= \phi x_{t-1} + \varepsilon_t^x \\r_{t+1} &= x_t + \varepsilon_{t+1}^r.\end{aligned}$$

x_t is the expected return, then r_{t+1} adds unexpected noise ε_{t+1}^r . We’re going to build a little model economy. People in the economy know the expected return x_t , but we don’t see it. So, we’ll figure out what prices they will set, knowing x_t ; then we’ll see if prices (that we can see) will forecast returns.

- (a) Using parameters $\phi = 0.94$, $\sigma(\varepsilon^x) = 0.0171$, $\sigma(\varepsilon^r) = 0.173$, start at $x_0 = 0$, plot 100 simulated years of r_{t+1} together with x_t by picking ε_t^x and ε_t^r from a random number generator. You should see how expected returns move slowly through time and how actual returns then move randomly around the expected return. The simulated return series should look reasonable, pretty much like actual annual returns you plotted in problem set 1, except for luck of which random numbers got picked. (I reverse-engineered these assumptions to generate $\sigma(E_t(r_{t+1})) = \sigma(x_t) = 5\%$, and $\sigma(r_{t+1}) = 18\%$.)
- (b) Suppose dividend growth is completely unpredictable. It follows

$$\Delta d_{t+1} = \varepsilon_{t+1}^d$$

using $\sigma(\varepsilon^d) = 0.10$ simulate this as well and plot expected and actual dividend growth. Would it be easy to tell this series apart from the return series just by looking at it?

- (c) Using the present value identity

$$d_t - p_t = E_t \sum_{j=1}^{\infty} \rho^{j-1} (r_{t+j} - \Delta d_{t+j})$$

find the dividend yield in terms of things we have already specified at time t , like $x_t, r_t, \Delta d_t$, etc. Hint: $E_t(r_{t+1}) = x_t$. $E_t(r_{t+2}) = \phi x_t$. Plot the dividend yield in your simulation as well. Use $\rho = 0.96$. Does the plot look like the actual log dividend yield – is it slow-moving like real d-p or fast-moving like returns?

Stop and smell the flowers: We wrote down a time series model for discount rates r , and dividend growth rates Δd . Now we are deriving what prices *should* be from the present value formula. We’re figuring out what the variables we can see ($dp, r, \Delta d$) would look like if this is how expected returns x_t , which we can’t see, and dividend growth evolved. You should also see how the price/dividend ratio (which we can see) *reveals* to us information about the true expected return x_t which we can’t see directly.

- (d) Now, we want to know, what would return forecasts look like if the world worked this way? First, find the return forecast analytically – express the model you have so far as

$$r_{t+1} = (\text{constant}) + (\text{coefficient}) \times (dp_t) + (\text{error}_{t+1})$$

and find the coefficient both as an algebraic expression and as a number. What does your model predict for the other regressions as well,

$$\begin{aligned} \Delta d_{t+1} &= (\text{constant}) + (\text{coefficient}) \times (dp_t) + (\text{error}_{t+1}) \\ dp_{t+1} &= (\text{constant}) + (\text{coefficient}) \times (dp_t) + (\text{error}_{t+1}) \end{aligned}$$

(Hint: The numbers should look a lot like the ones in real data.)

- (e) To check, run the return forecasting regression $r_{t+1} = a + b_r (d_t - p_t) + \varepsilon_{t+1}^r$ in your simulated data. Use 10,000 simulated data points to avoid luck of the draw. Do you get the same coefficient as in d? Does the R^2 look about the same as in real data? How about t statistics? Also find the R^2 for 5 year return forecasts. Why? You’re answering the question “suppose the data were generated by this model, what would the 5-year R^2 be?” In part d, you were able to find the model-predicted regression coefficient by algebra, but finding the predicted R^2 would be a huge pain. You learn here how easy it is to find answers to questions like that by simulation.

More explanation: This is your first simulation model. The idea is, “suppose the world worked this way – for example, suppose expected returns varied slowly through time and expected dividend growth did not, and prices are set by the present value relation. What would return, dp, and dividend growth data from that world look like?”

The key issue is that we cannot directly see x in the real world, so we have to evaluate this view of the worlds by observations of $r, \Delta d, dp$ alone. That’s why we have to substitute out for x and run regressions on $d - p$ rather than just look at x . This is a model with a “structural” version given by x , but a (usually) weaker set of “observable implications” given by the $\{r, \Delta d, d - p\}$. Our job in finance, like in science, is to tease out the “structure” from less-direct “observable implications”

Hints:

- (a) The model in part a is very important, and you’ll see it often in your finance life. All fixed-income models are based on it. I picked the parameters to generate $\sigma(r) = 0.18$. To do the simulation use the random number generator randn to pick random values for the ε_t , and then find the resulting x_t and r_t .
- (b)
- (c) In case you forgot, $\sum_{j=0}^{\infty} z^j = 1/(1 - z)$ for $|z| < 1$ will be useful. All these formulas apply to demeaned log variables, so the levels and sizes of numbers will be appropriate to those units. (If $D/P = 0.04$, $d - p = \log(0.04) = -3.2$.) The results depend on what your random number

generator picks. I suggest using `randn('state',number)` to start the random number generator. Then look at several different numbers, and try to pick one simulation that captures the typical result over many that you've looked at. There will be pathological simulations that give the “wrong” answer. This question *does not require lots of algebra*. If you're doing more than 2 or 3 lines, you're headed down the wrong path.

To produce a 100×1 matrix of independent random normals, use

```
randn(100,1)
```

To seed the random number generator so it produces the same string of numbers each time you run the program, use

```
randn('state',number)
```

To generate a random normal with mean μ and variance σ^2 , use

```
mu + sigma*randn(1)
```

to generate a simulation from an AR(1),

```
x = zeros(10,1);
x(1) = sigma*randn(1);
for t=2:10;
    x(t) = 0.94*x(t-1)+sigma*randn(1);
end;
```

Matlab hints

When you just need regression coefficients (not F , t , R^2 and so forth) you can get them quickly with

```
b = x\y
```

If you want a constant in the regression (yes) then the first column of x should have a constant in it, i.e.

```
b = [ones(size(x,1),1) x]\y
```

This cool command can run many regressions at the same time! For example, if

```
y = [r dp];
```

then

```
b = x\y;
```

gives you both the regression of r on x and the regression of dp on x at the same time! (The first column of the matrix b is the r regression coefficients, and the second column is the dp regression coefficients) This is useful now, and will be really useful to run all 25 Fama French portfolios on things in one line.

If you want the regression errors, then just go

```
e = y-b*x
```

In fact, you can program up your own OLS routine this way, and the formulas look just like the textbook (see `olsgmm.m`)

To do the same thing over and over again, use a "for loop." For example, you could generate the AR(1) impulse response function like this:

```
x = zeros(10,1);
x(1) = 1;
for t=2:10;
    x(t) = 0.9*x(t-1);
end;
```

The program then repeats the statements inside the `for..end` block replacing i by 2,3,...10 in turn.

Be careful not to use the first element of the dividend growth. We don't have that data, so it is coded as NaN (not a number). If you see NaN results, it means you're unintentionally using it.

The identity (1) does not include means. It is meant to be applied to series that have had their means removed, or to apply only to b coefficients and variances that don't care about means.