Business 35150 **John H. Cochrane**

Problem set 9 answers

1. -0.05, 0.05, 0.95, 1.05

2.

$$\begin{aligned} -p_t^{(n)} &= \left(0 - p_t^{(1)}\right) + (p_t^{(1)} - p_t^{(2)}) + ... + (p_t^{(n-1)} - p_t^{(n)}) \\ p_t^{(n)} &= -\left(y_t^{(1)} + f_t^{(2)} + ... + f_t^{(n)}\right) \end{aligned}$$

3.

$$y_t^{(n)} = \frac{1}{n}\left(y_t^{(1)} + f_t^{(2)} + ... + f_t^{(n)}\right)$$

4.

$$y_t^{(n)} = \frac{1}{n}\left(y_t^{(1)} + E_t\left(y_{t+1}^{(1)}\right) + ... + E_t y_{t+n-1}^{(1)}\right)$$

5.

$$\begin{aligned} -p_t^{(n)} &= \left(0 - p_{t+n-1}^{(1)}\right) + \left(p_{t+n-1}^{(1)} - p_{t+n-2}^{(2)}\right) + \left(p_{t+n-2}^{(2)} - p_{t+n-3}^{(3)}\right) + ... + \left(p_{t+1}^{(n-1)} - p_t^{(n)}\right) \\ -p_t^{(n)} &= r_{t+n}^{(1)} + r_{t+n-1}^{(2)} + r_{t+n-2}^{(3)} + ... + r_{t+1}^{(n)} \end{aligned}$$

It should bother you! But bond returns are negatively serially correlated. If r_{t+1} is high r_{t+2} must be lower, as bonds always end up at zero.

6.

$$\begin{aligned} f_t^{(2)} &= p_t^{(1)} - p_t^{(2)} = -p_{t+1}^{(1)} + p_t^{(1)} + p_{t+1}^{(1)} - p_t^{(2)} \\ f_t^{(2)} &= y_{t+1}^{(1)} + \left(r_{t+1}^{(2)} - y_t^{(1)}\right) \\ f_t^{(2)} - y_t^{(1)} &= \left(y_{t+1}^{(1)} - y_t^{(1)}\right) + \left(r_{t+1}^{(2)} - y_t^{(1)}\right) \end{aligned}$$

7.

$$\begin{aligned} f_t^{(2)} - y_t^{(1)} &= \left(y_{t+1}^{(1)} - y_t^{(1)}\right) + \left(r_{t+1}^{(2)} - y_t^{(1)}\right) \\ f_t^{(2)} - y_t^{(1)} &= a_y + b_y\left(f_t^{(2)} - y_t^{(1)}\right) + \varepsilon_{t+1}^y + a_r + b_r\left(f_t^{(2)} - y_t^{(1)}\right) + \varepsilon_{t+1}^r \\ 0 &= a_y + a_r \\ 1 &= b_y + b_r \\ 0 &= \varepsilon_{t+1}^y + \varepsilon_{t+1}^r \end{aligned}$$

we have complementary regressions again. A high forward spot spread must signal an increase in yield or a high expected return, just as a high dividend yield must signal an increase in return or a low dividend growth.

Part II

The bootstrap and monte carlo. Start with the regression in real data

```
sample
   19461231    20121231
Regression of log  returns on log  D/P 1926-today
                            b      se(b)        R2
1 Yr. return            0.122      0.047     0.097
1 Yr. dp                0.949      0.039     0.901
```

This should be familiar by now.

```
correlation matrix of errors
    1.0000   -0.5942
   -0.5942    1.0000
```

The errors are strongly negatively correlated. This is a deep point – return shocks are likely to come from price shocks with no change in dividends. For our point, you do not want to simulate assuming the errors are independent of each other.

```
                                br       bdp
    assumed coefficients    0.1221    0.9486
mean of sim coefficients    0.1678    0.8839
std errors from regression  0.0465    0.0393
Bootstrap se                0.0740    0.0673
pct sim > observed         71.7400   14.7700

monte carlo
mean of sim coefficients    0.1681    0.8833
monte carlo se              0.0753    0.0688
percent of sim > observed  71.7000   14.4400

bootstrap with null = 0
assumed coefficients        0.0000    0.9400
mean of sim coeff           0.0516    0.8693
monte carlo se              0.0817    0.0744
percent of sim > observed  17.4800   10.1000
```

Start with the bootstrap. The regression is biased up! The true br is 0.12, but the mean in the bootstrap is 0.17. The bdp is biased down, and these two turn out to be related observations. br is biased up because bdp is biased down, and the errors are correlated. (AR(1) with large coefficients are biased down in small samples.) The return standard error from the regression is small by almost half! So, the mighty $0.122/0.0465 = 2.62$ t statistic is in fact only a $0.122/0.074 = 1.65$ t statistic. The dp standard error is also too small by half. Regression standard errors can be way wrong. Since the estimate is biased, you will not be surprised that the chance of seeing a larger br is more than 50%, in fact it's 72%, and similarly for bdp.

How much of this is non-normal errors? The monte carlo uses normal errors... and finds almost exactly the same numbers. Monte carlo vs. bootstrap does not really make much difference for this application.

Now, recenter the null at br=0, bdp=0.94. You still see the upward bias of br and downward bias of bdp – if it's truly zero we still see 0.056 on average, and (these are related) 0.86 for bdp. 17.5% of simulated observations have a br value greater than the 0.12 we see in the data. This means that, using these correct statistics, we do not reject the null hypothesis that br = 0 at the 10% level.

All of this seems really depressing. How did I spend so much time on an "insignificant" regression? Well, as I hope you know by now, the point estimates are economically very significant. But when you look at the *joint* plot of return and dp, you see how unlikely it is to see *both* a high return coefficient *and* a high dp coefficient. This is really the key. All of the simulations with a spuriously high (>0.12, when the true value is 0) return coefficient *also* have a spuriously low (<0.94) dp coefficient. My "dog that didn't bark" paper works out this statistical point. When you take both return and dp observations together, yes, it is highly significant.

Here are the plots:

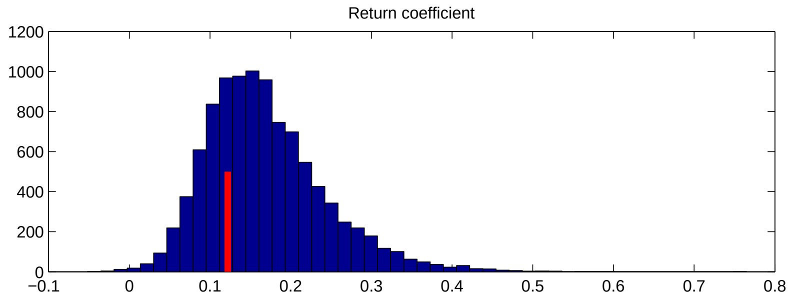

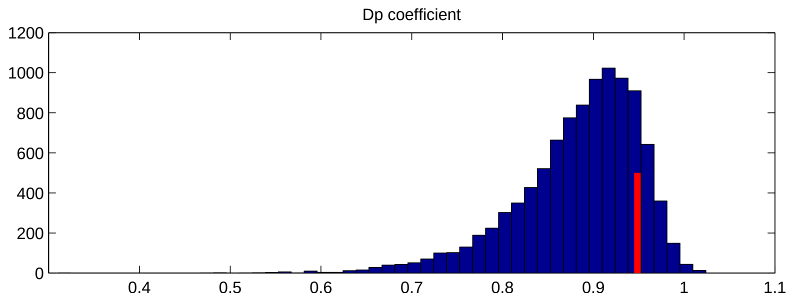

The red lines are the "true" or "null" values, as estimated from our regression. As you can see, the return coefficient is biased up, and the dp coefficient is biased down. The distributions are far from normal – 5% probability values based on the standard errors would be way off the actual probability values. The joint distribution shows an interesting fact – samples with high return coefficients have low dp coefficients.

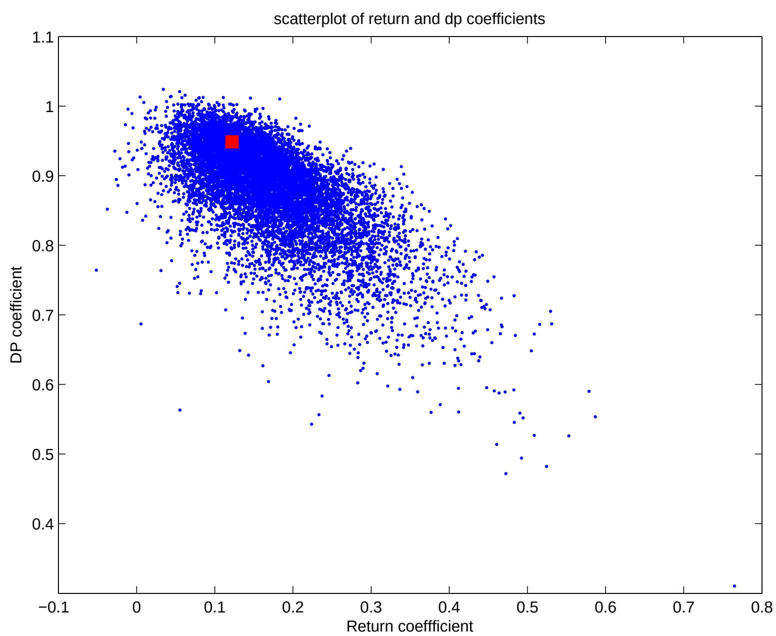

The monte carlo graphs look the same. The non-normal *sampling distributions* do not come from non-normal errors.

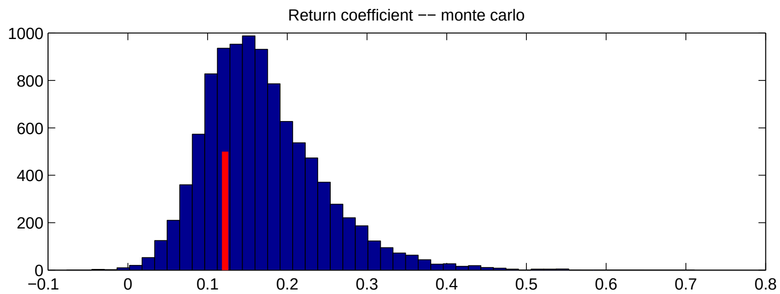

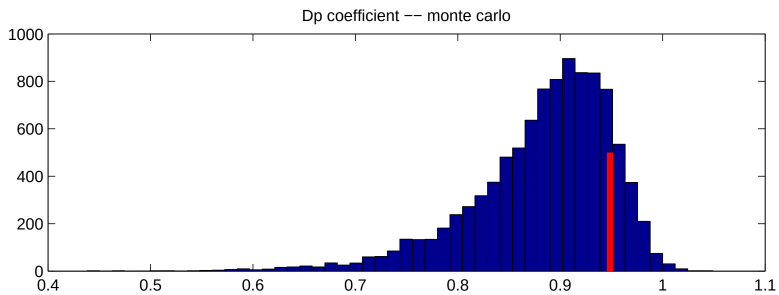

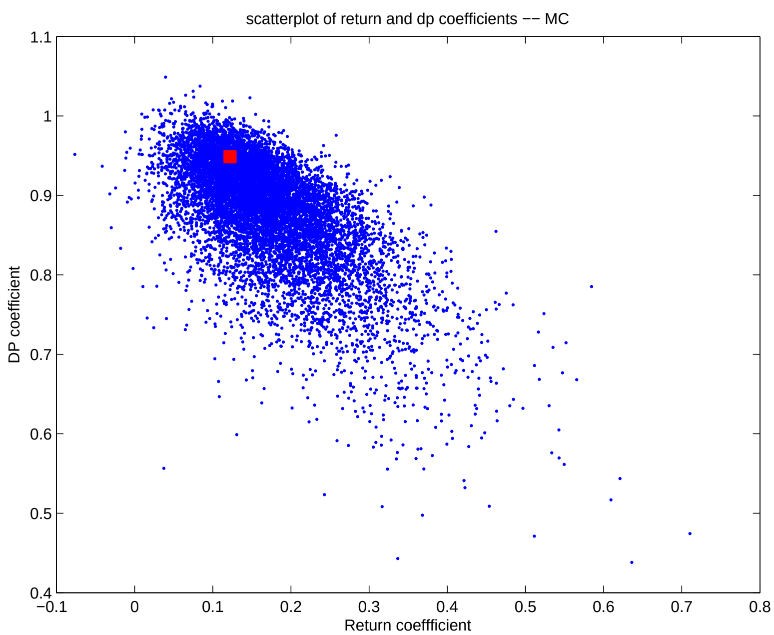

The graphs based on the null of $b_r = 0$ show that the return distribution is biased up, not so much because the mode is biased up but because of the fat right tail. You can see quite a bit of probability mass above our sample value of 0.12. Correspondingly, the dp coefficient is biased down.

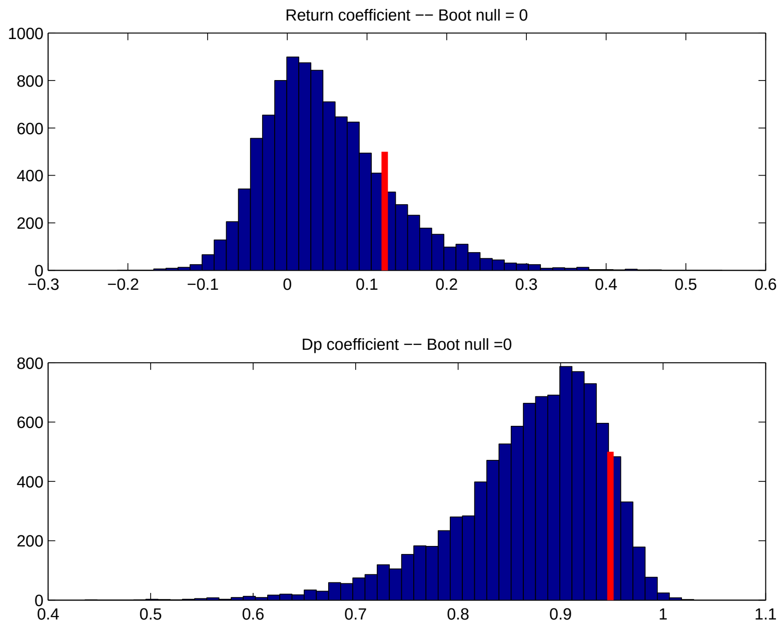

Here the *joint* distribution of return and dp matters. Since samples that have a too high b_r *also* have too low b_{dp}, the chance of seeing *both* high b_r and b_{dp} is very low – our data is at the edge of the galaxy.

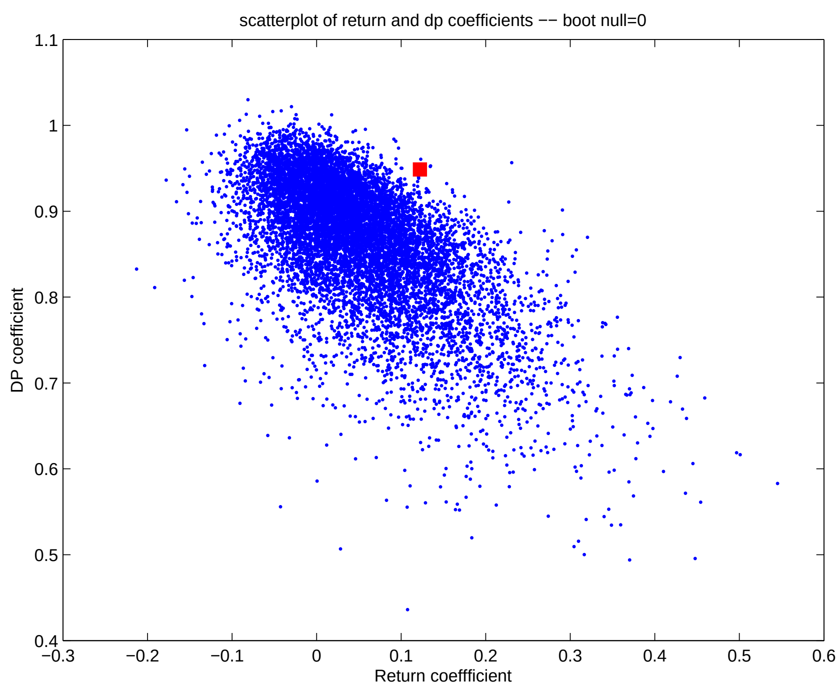

Figure 1: