Problem Set 2 Answers

Part I. Reading questions:

1. Based on the regressions shown in the readings, $R_{t+1} = a + 4 \times DP_t + \varepsilon_{t+1}$, so expected returns change by four times as much as the DP ratio, so the expected return rises 4 percentage points

2.

   (a) Expected return (average return is ok) $E(R_{t+1}) = 0$

   (b) Conditionally expected return, time-t expected return. $E_t(R_{t+1}) = 1$

   (c) Conditional standard deviation of return $\sigma_t(R_{t+1}) = 0.4$

   (d) Standard deviation of expected returns $\sigma\left[E_t(R_{t+1})\right] = \sigma(a + bx_t) = 1 \times \sigma(x_t) = 0.3$

   (e) Standard deviation of return $\sigma(R_{t+1}) = \sqrt{\sigma^2(a + bx_t) + \sigma^2(\varepsilon)} = \sqrt{0.3^2 + 0.4^2} = 0.5$

   (f) R squared. $R^2 = \sigma^2(a + bx_t)/\sigma^2(R_{t+1}) = (0.3/0.5)^2 = (3/5)^2 = 0.6^2 = 0.36$

3. If we drive prices up, for given cash flows, then returns will be lower. Higher fear drives prices down, and this is how the greater returns are created. This is an important and common confusion. One way to keep this straight is to think of "required returns" rather than "expected returns."

4. If these are point estimates, we'd conclude there is a bug in the program. If these are someone's opinions or a test, we'd conclude the author was confused. $r_{t+1} = (d_{t+1} - d_t) - (p_t - d_t)$ means $b_r = b_d + 1$.

5. The answer is that the dividend and return terms must be perfectly correlated. Any surprise increase in returns must come from a surprise increase in dividends. This is pretty obvious in the one period case

$$R_{t+1} = \frac{D_{t+1}}{P_t}$$

means

$$\begin{aligned} P_t &= \frac{D_{t+1}}{R_{t+1}} \\ p_t &= d_{t+1} - r_{t+1} \end{aligned}$$

We used the expected value implication

$$p_t = E_t d_{t+1} - E_t r_{t+1}$$

but it also means the unexpected values follow the same identity

$$0 = (d_{t+1} - E_t d_{t+1}) - (r_{t+1} - E_t r_{t+1})$$

$$d_{t+1} - E_t d_{t+1} = r_{t+1} - E_t r_{t+1}$$

In statistical terms, the unexpected parts of dividends and returns are perfectly correlated. In intuitive terms, once you pay the price, the *way* you get an unexpected returns is by getting an unexpected dividend. Well, duh. Hence, returns and dividends, given the price, are perfectly correlated and their log *difference* is known at time t even though each element (r and d separately) are not known at time t.

In the present value version,

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

means, yes,

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

but also

$$0 \approx (E_t - E_{t-1}) \sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j} - (E_t - E_{t-1}) \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j}.$$

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

so long-run expost dividends and returns are perfectly correlated. Duh again, *the only way to get a good long run return is to get higher than expected dividends.* We spend so much time thinking about returns as caused by price rises, that this is a vital thought check. In the end, it's all about dividends.

Part II Computer questions

My answers are based on the sample 19261231 20121230

1. Here is my plot of actual and approximate log dividend growth. As you can see, the approximation is excellent. The other identities come out very closely if you use real rather than approximate dividend growth.

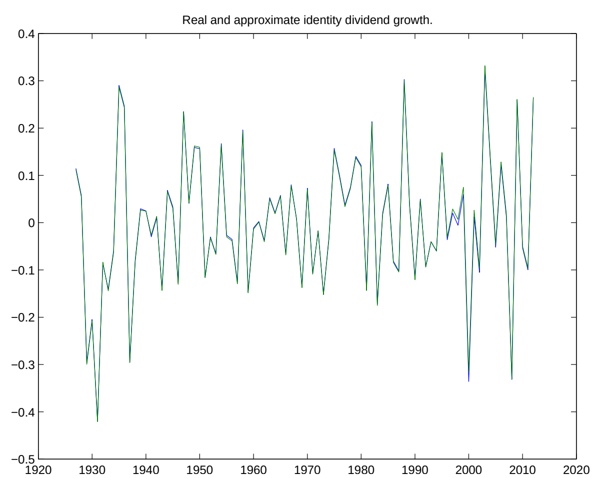

2. Regression

   (a) My coefficients

b_r	b_d	b_{dp}
0.0966	-0.0048	0.9361

(b) The identity:

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

if you run both sides on dp_t, you get

$$b_r = 1 - \rho b_{dp} + b_d$$

More formally, substitute the regression equations (leaving out constants) into the identity

$$\begin{aligned} r_{t+1} &= -\rho dp_{t+1} + dp_t + \Delta d_{t+1} \\ b_r dp_t + \varepsilon^r_{t+1} &= -\rho \left(b_{dp} dp_t + \varepsilon^{dp}_{t+1} \right) + dp_t + \left(b_d dp_t + \varepsilon^d_{t+1} \right) \end{aligned}$$

This equation must hold for every value of dp_t so the terms multiplying dp_t and the other terms must separately be equal

$$\begin{aligned} b_r dp_t &= -\rho \left(b_{dp} dp_t \right) + 1 dp_t + b_d dp_t \\ b_r &= -\rho b_{dp} + 1 + b_d \end{aligned}$$

$$\varepsilon^r_{t+1} = -\rho \varepsilon^{dp}_{t+1} + \varepsilon^d_{t+1}$$

The two sides of the coefficient identity:

b_r	$1 - \rho b_{dp} + b_d$
0.0966	0.0966

They are *exactly* the same. That is, of course because I used the identity to construct the dividend growth series. If you use the real dividend growth series, the identity is only approximate.

(c) If you take residuals of both sides of the identity ($\varepsilon^x_{t+1} = x_{t+1} - a_x - b_x dp_t$) you get

$$\varepsilon^r_{t+1} = -\rho \varepsilon^{dp}_{t+1} + \varepsilon^d_{t+1}.$$

Or follow the more formal derivation above. The two sides of that identity (first 5 observations)

0.1598	0.1598
0.2089	0.2089
-0.2655	-0.2655
-0.4377	-0.4377
-0.7059	-0.7059

these are *exactly* the same

(d) My coefficients

b_r	b_d	b_{dp}
0.0966	0.0027	0.9361

and the identity

b_r	$1 - \rho b_{dp} + b_d$
0.0966	0.1040

as you can see they're very close.

(e) The identity

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

says (duh) that the *return* is composed of dividend yield change and dividend growth. When we run a regression, we decompose the return into parts predictable from dp and residuals. So the predictable part of returns is composed of the predictable part of dividend yield and dividend growth, and the residual unpredictable part of return is also composed of unpredictable parts of dividend yield and dividend growth.

3. The impulse response functions. First, the dp shock

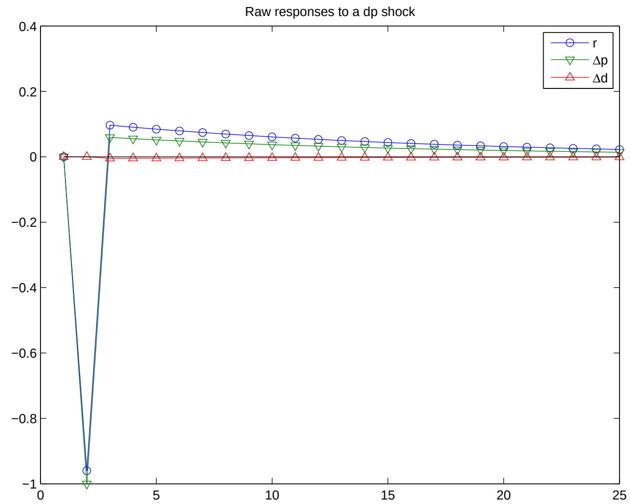

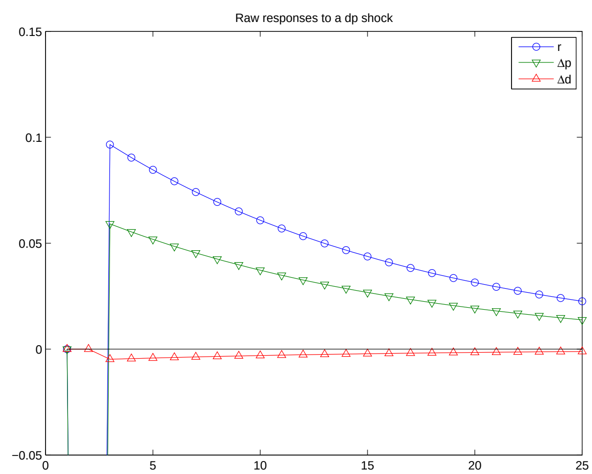

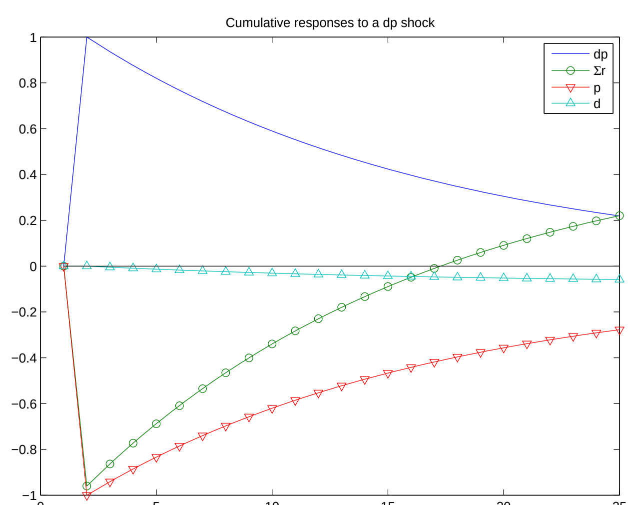

This looks pretty much as in class. Since the b_d coefficient is nearly exactly zero, the response of dividend growth to the dp shock is basically zero now as well. Second, the dd shock,

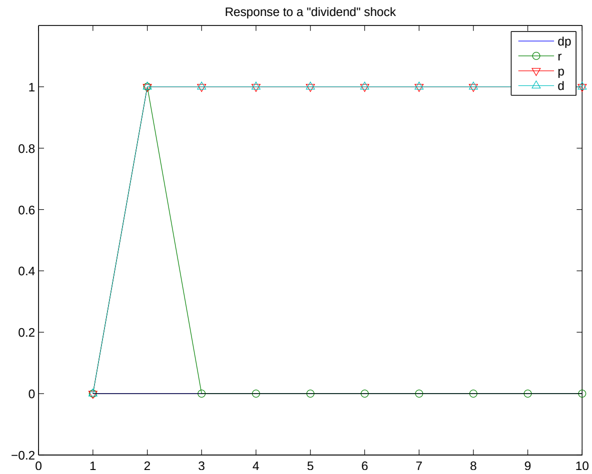

As before, this is essentially a pure random walk.

Why the weird difference between the impact and subsequent returns? When we raise *expected returns* $E_t r_{t+j}$ in the future, this lowers *today's* price, so we get a sharp negative impact return, then the higher expected returns.

4. The 1970s view is that

   (a) $b_r = 0$. $b_r = 0$ means $b_d = \rho b_{dp} - 1$

b_r	$b_d = \rho b_{dp} - 1$	b_{dp}
0	-0.0966	0.9361

   If we want $b_r = 0$, we need to assume $b_d \approx -0.1$. I wrote a whole paper on this titled "The dog that didn't bark." Interpretation: now a higher *dp* means a lower price for given dividend. With no change in expected returns, this means that a low price (high dp) must signal lower future dividends. No, both returns and dividend growth cannot fail to be forecastable, or dp would never change in the first place. Our 1970s professor would brush his long hair back and answer the critic: "Prices are high because people expect higher dividend growth in the future." If people are not systematically wrong in their expectations, then times of higher price/dividend ratio will in fact be followed on average by higher dividend growth.

   (b) In sum, we are now looking at what the hypothetical system

$$
\begin{aligned}
r_{t+1} &= 0 \times dp_t + \varepsilon^r_{t+1} \\
\Delta d_{t+1} &= -0.0966 \times dp_t + \varepsilon^d_{t+1} \\
dp_{t+1} &= 0.9361 \times dp_t + \varepsilon^{dp}_{t+1}
\end{aligned}
$$

   looks like when fed our two shocks. Here are the new impulse response functions

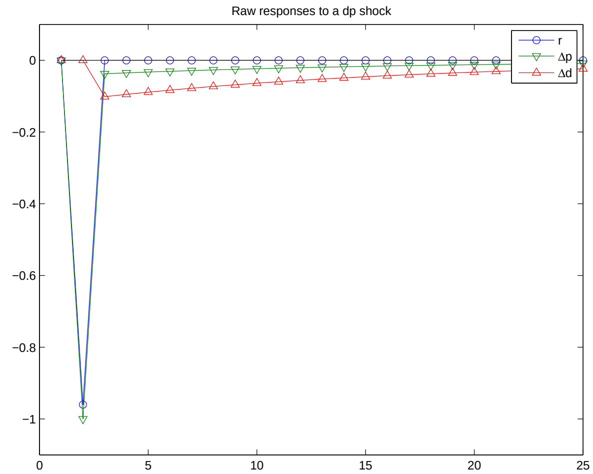

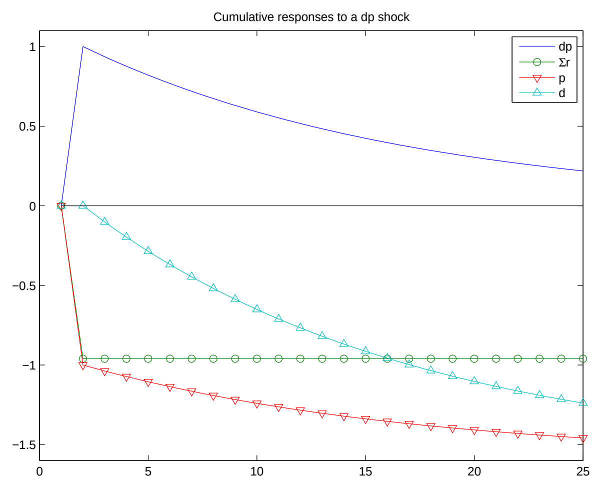

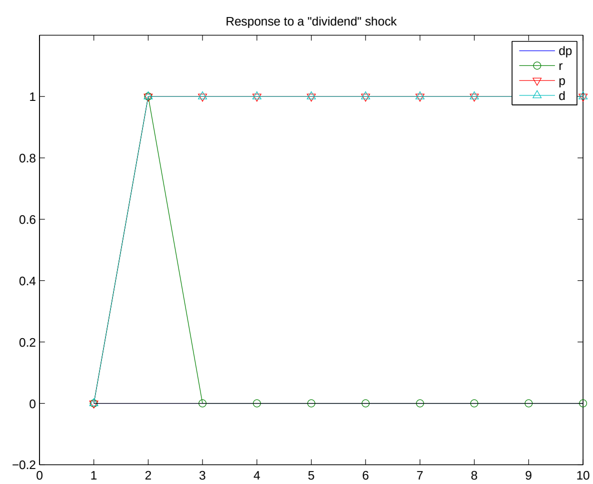

(c) The dividend growth shock as we have defined it looks the same. As before, it's a change in dividend growth with no change in dividend yield. If dividend yield doesn't change, nothing on the right side changes for $t = 2$, so any effect is purely transitory. It does not move expected returns for times past the shock at all. The return with the shock is positive, as you got more dividend and the price rose exactly as much, keeping the dp ratio constant. Once again, it's a pure "permanent cashflow" shock.

The dividend yield shock gives a really different response. What's going on? Well, the dividend/price ratio goes up with no change in dividend growth, which still means a big price decline and big negative return in the shock period by addition. Now it means *no change in expected returns – the higher dp has no effect on expected returns.* But now it implies that *future* dividend growth must be lower. And it is – notice the red line, the lower dividend growth following the shock (top graph).

(d) So now the names of the two shocks are different. I'd call the dividend yield shock a "shock to news about future dividend growth with no change in current dividend growth" or something similar. Dividend growth is not iid anymore, it has a permanent and transitory component and the two shocks separate them. *To account for variation in pd ratios, with no variation in expected return, you have to believe there is variation in future dividend growth beyond what you see in current dividend growth, so that higher price/dividend ratios correspond to higher future long-run dividend growth.*

5. The question is

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

(a) and

(b) Here's are the plots. In the first plot, the red line is the x, expected return series, and the black line is the r, or actual return series. Big points:

i. You should see here how the model works. You see the hidden expected return x_t series wanders slowly through time, and then the actual return r_t adds some noise to that.

ii. You should also see how hard it would be to detect the varying mean return if all you could not see the red lines, because the black lines we see are so volatile. *The large volatility of actual returns makes it very hard to see any variation in expected returns.*

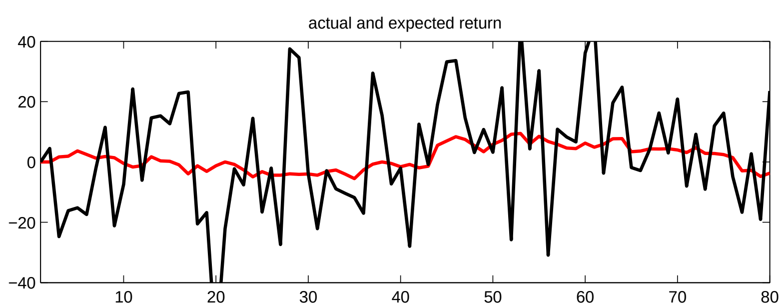

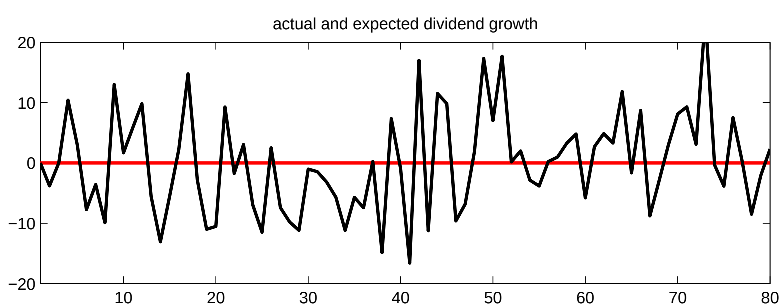

c) The dividend yield is

$$
\begin{aligned}
d_t - p_t &= E_t \sum_{j=1}^{\infty} \rho^{j-1} \left(r_{t+j} - \Delta d_{t+j} \right) \\
&= x_t + \rho\phi x_t + \rho^2 \phi^2 x_t .. \\
&= \frac{1}{1-\rho\phi} x_t .
\end{aligned}
$$

See the plot. It's supposed to give you confidence that the model is pretty reasonable; this looks a lot like the real dividend yield. You can also see how we've captured the very long swings in DP that are true in our world.

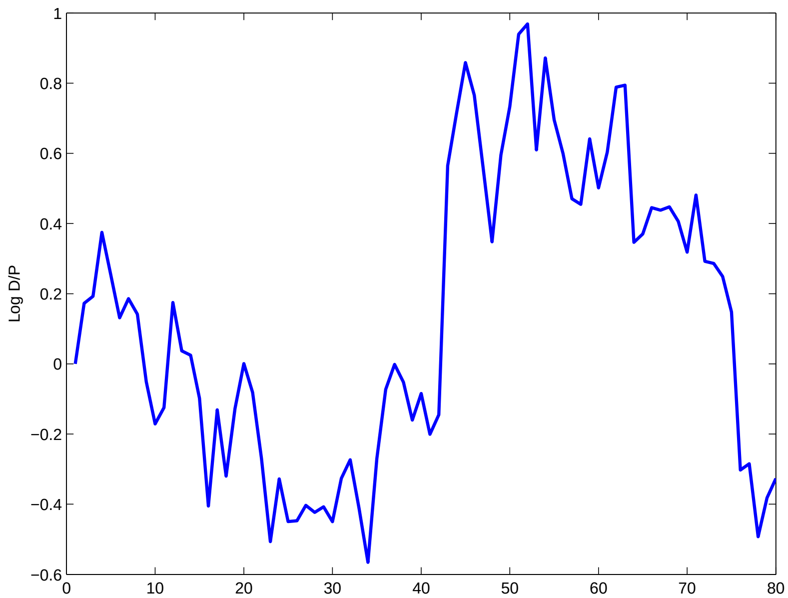

d-p simulation from a world with time-varying expected return

c) We have

$$
\begin{aligned}
d_t - p_t &= \frac{1}{1-\rho\phi} x_t \\
r_{t+1} &= x_t + \varepsilon^r_{t+1}
\end{aligned}
$$

Substituting,

$$r_{t+1} = (1-\rho\phi)(d_t - p_t) + \varepsilon^r_{t+1}$$

so we expect a coefficient $b_r = 1 - \rho\phi = 1 - 0.94 \times 0.96 = 0.098$. Of course, we expect $b_d = 0$ and $\phi = \phi = 0.94$.

d) Here are my regressions. From the return world

```
simulation using expected returns
assumed sigma(epsilon_x), sigma(epsilon_r), sigma(epsilon_d)
    0.0171    0.1729    0.1000
Regression of log returns on lagged d-p in simulated data with  10000 observations
                          b          se        t(b)         R2
1 Yr.                  0.093       0.003      27.685       0.071
5 Yr. no overlap       0.395       0.017      22.733       0.206
```

All I asked for in the problem set is the coefficient. We expected 0.098, we got 0.093 with a standard error of 0.003, well in the right range. (A longer sample would get closer to 0.098.) The other numbers just fill in to show we're getting very similar numbers to what we got in the data last week. We see the expected R^2 of 7%, and b and R^2 rise steadily with horizon.

This is a selling point for simulations. We were able to answer "if this is how the world works, what do you expect for b_r?" analytically. However, figuring out this model's predictions for long-horizon returns and R^2 looked to be an algebraic nightmare. Luckily, we can just let the computer work it out for us.

The t statistics are ridiculous, of course, because we have a sample size of 10,000. To see t stats like those in the data, we would have to use a sample size similar to that in the data. But then the coefficient would have suffered some luck of the draw. The right way to evaluate small-sample performance is with a Monte Carlo which we will do later. The standard errors are only useful here to give a sense of whether the simulation is long enough.

We know that the dividend growth coefficient b_d should be zero. We can check that in the simulation too:

```
Regression of log dividend growth on lagged d-p in simulated data with  10000 observations
                              b          se       t(b)         R2
1 Yr.                    -0.000       0.002     -0.022      0.000
5 Yr. no overlap         -0.012       0.010     -1.243      0.001
```

Note: This simulation model has a severe shortcoming, that I assumed the return ε^r and expected-return (ε^x, hence ε^{dp}) shocks are uncorrelated. In the real world, these shocks are highly correlated – a rise in expected returns sends prices down, and hence is negatively correlated with a rise in returns. We didn't use that feature, and I neatly avoided any calculations in which this would cause trouble. (The regression of returns on past returns would have come out wrong.) I just want to warn you not to use this model for other purposes without getting the shock correlations right.

The right way to set up this model is a bit more complex. Write

$$\begin{aligned} x_t &= E_t r_{t+1} \\ x_t &= \phi x_{t-1} + \varepsilon^x_t \\ \Delta d_{t+1} &= \varepsilon^d_{t+1} \end{aligned}$$

with ε^x and ε^d uncorrelated. Then find

$$dp_t = E_t \sum_{j=1}^{\infty} \rho^{j-1} \left(r_{t+j} - \Delta d_{t+j} \right) = \frac{x_t}{1-\rho\phi}$$

finally find

$$\begin{aligned} r_{t+1} &= -\rho dp_{t+1} + dp_t + \Delta d_{t+1} \\ &= -\frac{\rho x_{t+1}}{1-\rho\phi} + \frac{x_t}{1-\rho\phi} + \varepsilon^d_{t+1} \\ &= -\frac{\rho\left(\phi x_t + \varepsilon^x_{t+1}\right)}{1-\rho\phi} + \frac{x_t}{1-\rho\phi} + \varepsilon^d_{t+1} \\ r_{t+1} &= x_t + \left(\varepsilon^d_{t+1} - \frac{\rho}{1-\rho\phi}\varepsilon^x_{t+1} \right) \end{aligned}$$

Now you see returns are of the same form as before, x_t is the expected return and there is a shock, but the return shock here will have a strong negative correlation with the expected return = dividend yield shock.