TAKE-HOME FINAL EXAM

(1) We consider three variants of a simple New Keynesian model. They share "IS" and "Phillips Curve" equations:

$$\text{IS}: \qquad r_t = E_t\pi_{t+1} + \beta E_t[y_{t+1} - y_t] + \varepsilon_{2t} \qquad (1)$$

$$\text{PC}: \qquad \pi_t = \nu E_t\pi_{t+1} + \theta y_t + \varepsilon_{3t}. \qquad (2)$$

They differ in their monetary policy reaction functions, which we label b, c, and f for "backward", "contemporaneous", and "forward":

$$\text{f}: \qquad r_t = \alpha E_t\pi_{t+1} + \delta y_t + \varepsilon_{1t} \qquad (3)$$

$$\text{c}: \qquad r_t = \alpha\pi_t + \delta y_t + \varepsilon_{1t} \qquad (4)$$

$$\text{b}: \qquad r_t = \alpha\pi_{t-1} + \delta y_{t-1} + \varepsilon_{1t}. \qquad (5)$$

When handed off to gensys, these models return solutions in the form

$$y_t = Gy_{t-1} + \Theta_y\left(\sum_{s=0}^{\infty}\Theta_m^s\Theta_z E_t z_{t+s+1}\right) + Hz_t, \qquad (6)$$

where y is the vector of variables (here π, r, y, in that order) and z is the vector of exogenous disturbances. Note that because the program makes all equations have the same most-advanced time index, the z_t vector is $[\varepsilon_{1t}, \varepsilon_{2,t-1}, \varepsilon_{3,t-1}]'$ in the case of models b and c, and $z_t = [\varepsilon_{1,t-1}, \varepsilon_{2,t-1}, \varepsilon_{3,t-1}]'$ in the case of the f model. These disturbances are assumed to be i.i.d. with zero mean, but because of this timing $E_t z_{t+1}$ is non-zero and the Θ parts of the solution therefore matter. Because of the i.i.d. assumption, only $\Theta_y\Theta_z$ matters, however.

The solutions for the three cases for a particular set of numerical values of the parameters are

f: $G = 0$ $H = 0$

$$\Theta_y\Theta_z = \begin{bmatrix} 0.3571429 & -0.35714286 & -1.000000e+00 \\ -1.0714286 & 0.07142857 & 0 \\ 0.7142857 & -0.71428571 & 0 \end{bmatrix}$$

c: $G = 0$ $H = \begin{bmatrix} -0.2325581 & 0 & 0 \\ 0.6976744 & 0 & 0 \\ -0.4651163 & 0 & 0 \end{bmatrix}$

$$\Theta_y\Theta_z = \begin{bmatrix} -0.3547864 & -0.2325581 & -0.7441860 \\ -0.4240130 & -0.3023256 & -0.7674419 \\ -0.3374797 & -0.4651163 & 0.5116279 \end{bmatrix}$$

b: $\quad G = \begin{bmatrix} -0.2233596 & 0 & -0.02030542 \\ 1.1000000 & 0 & 0.10000000 \\ -0.5442236 & 0 & -0.04947487 \end{bmatrix}$ $\qquad H = \begin{bmatrix} -0.2030542 & 0 & 0 \\ 1.0000000 & 0 & 0 \\ -0.4947487 & 0 & 0 \end{bmatrix}$

$$\Theta_y \Theta_z = \begin{bmatrix} -3.006442e-01 & -2.030542e-01 & -6.692914e-01 \\ 0 & 0 & 0 \\ -3.967338e-01 & -4.947487e-01 & 4.364071e-01 \end{bmatrix}$$

For each of the three models, assuming that these numbers describe the true model generating data we will observe, answer these questions:

(a) Is the model invertible — that is, from observations on current and past values of the three variables π, r, y, can one recover the structural disturbance vector of ε's?

There are two ways in principle that invertibility could fail in this setting. One is for the matrix of coefficients on the current structural shocks to fail to be full rank. The other is for there to be a moving average operator generated by the need to forecast the z_t's, and for this moving average operator to be non-invertible. In model f the z_t vector is lagged ε's, but since $H = 0$, there is no MA operator. Only current shocks matter, and $\Theta_y \Theta_z$ is non-singular. So the model is invertible. In models c and b the columns of H corresponding to the two lagged shocks are zero, so again there is no MA operator. In both cases the matrix of coefficients on the current shocks is the first column of H pasted in front of the last two columns of $\Theta_y \Theta_z$. The first column of $\Theta_y \Theta_z$ is coefficients on $E_t \varepsilon_{1,t+1}$, which is zero, so we don't use it. The matrix constructed this way is non-singular for both these models also.

(b) Are there zero restrictions on the contemporaneous coefficient matrix of an SVAR that would allow correct identification of the monetary policy reaction function?

Here are the Γ_0's, the contemporaneous coefficient matrices in the SVAR, for the three models:

f: $\qquad \begin{bmatrix} 0 & -1 & 0.1 \\ 0 & -1 & -1.5 \\ -1 & 0 & 0.5 \end{bmatrix}$

c: $\qquad \begin{bmatrix} -1.1 & 1 & -0.1 \\ 0 & -1 & -1.5 \\ -1 & 0 & 0.5 \end{bmatrix}$

b: $\qquad \begin{bmatrix} 0 & 1 & 0 \\ -1.0397 & -1 & -1.5945 \\ -1.1787 & 0 & .4838 \end{bmatrix}$

The matrix of coefficients on the current shocks, inverted, is the matrix of coefficients on current variables in the SVAR representation. In the f model that matrix has three zeros, enough to satisfy the order condition for identification. But the

zeros are placed so that they would not allow identification. In particular, the IS curve and the monetary reaction function, the first two equations, both have the same single zero restriction: inflation does not enter them contemporaneously. This restriction arises because the solution makes all the variables serially uncorrelated, so the $E_t\pi_{t+1}$ terms in the model are all identically zero. Obviously this restriction, though implied by this very simple model, is not one expects to hold in the actual data. Another possibly puzzling aspect of the f model is that I set it up with the r coefficient in the reaction function normalized to -1, rather than 1, so that a positive shock to the equation lowers r, raises π and raises y. Finally, a minor glitch was that I set the coefficient $\delta = -.1$, instead of $\delta = .1$ as in the other two models. Though unrealistic, this had only small effects on the coefficient matrices in the solution.

The c model has only two zero restrictions, and the policy reaction function is certainly not identified by those restrictions. The sum of the first two equations, for example, satisfies the zero restriction on the third equation, so an orthonormal transformation of the first two (the sum and difference of them) would yield one equation with no zeros (like the first) and one with a zero that satisfies the third equation's restriction.

The b model has three restrictions, and the two on the monetary policy reaction function obviously suffice to identify it. This is apparent from the original setup of the model: The one-step-ahead prediction error in r is the structural disturbance to the monetary policy equation, so the r equation in a reduced form VAR is the structural reaction function. In fact, this pattern of zeros delivers exact identification for all three equations, as was pointed out in the notes on SVAR identification for this course.

(c) Could the Clarida-Galí-Gertler strategy of using lagged values of π, y and r as instrumental variables work to allow estimation of the policy reaction function?

The CGG identification strategy rests on three assumptions in their model (which is like the f model here). They assume that lagged values of any variables in the system are in the information set at t, that the shocks are serially uncorrelated, and (though this assumption is not discussed explicitly by them) that the lagged instruments are correlated with the included variables. The first two assumptions are met in model f for this problem. However model f, because its solution implies $G = 0$ and no lagged disturbances enter the solution, implies that all three series are serially uncorrelated. Therefore the CGG strategy would not work. (In the original CGG work the rest of the model is not laid out, so we can't check this assumption. The tests of overidentifying restrictions that CGG display check the first two assumptions, not the third.)

In model b we do not need the CGG strategy, but if we wanted to use an instrumental variable for lagged π or y, this would work, as G is non-zero, so the lagged instrument would be correlated with the variable it is instrumenting for. However

using twice-lagged values as instruments for both π_{t-1} and y_{t-1} would not work. G in this model is of rank 1, meaning that if we tried to instrument for *both* rhs variables, the second stage of two-stage least squares would show perfect collinearity.

In model c we do have simultaneity, so it would be helpful if there were eligible lagged-variable instruments. But model c, like model f, has $G = 0$, so again the data are implied to be serially uncorrelated and lagged-variable instruments will not provide consistent estimates.

(2) An estimated AR model for US log GDP emerges as

$$y_t = 0.004732468 + .000004261858t + 1.371202y_{t-1} - 0.371252y_{t-2} + \varepsilon_t \qquad (7)$$

The variance of ε_t is estimated as 8.87316e-05. The two initial values of y, y_1 and y_2, are 7.359161 and 7.357973. Note that the variable "t" in the regression equation is 1 for the first observation, 2 for the second, etc, meaning that it has the value 3 in the first observation that was used in estimation.

Determine whether or not the two initial condition values are so far from the deterministic trend line for this model that the model implies such values are unlikely to occur again soon. Justify your conclusion.

It is a problem with this question that numerical answers depend on the sum of coefficients on lagged y's, and I gave you values that allowed this difference to be calculated to only two significant figures accuracy. Since the answer didn't depend on fine points of numerical accuracy, this was more or less OK, but the answers below are based on a higher order of accuracy.

The first step in answering this question is to find the trend line. Note that because the sum of coefficients on lagged y's is (just barely) less than one, the estimated model is indeed stationary around a trend line. We can find a trend line $\bar{y}_t = a + bt$ such that (7) can be rewritten as

$$y_t - \bar{y}_t = 1.371202(y_{t-1} - \bar{y}_{t-1}) - 0.371252(y_{t-2} - \bar{y}_{t-2}) + \varepsilon_t .$$

Some algebra shows that this implies that the trend line has coefficients $a = -995.9$, $b = .08602$. Since this implies that the trend growth rate is over 8% per quarter and the trend value of GDP (not log GDP) was close to zero at time 0, just before the start of the sample, we may already be suspicious. But to know how unusual the initial values are implied to be, we have to calculate the unconditional variance of y. This is not a completely trivial task for a second-order AR. A way to do it that we discussed in class is using a doubling algorithm. That is, first rewrite the model as a two-dimensional first-order model

$$\begin{bmatrix} \tilde{y}_t \\ \tilde{y}_{t-1} \end{bmatrix} = \begin{bmatrix} \rho_1 & \rho_2 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} y_{t-1} \\ y_{t-2} \end{bmatrix} + \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix} ,$$

where $\tilde{y}_t = y_t - \bar{y}_t$. Letting B be the right-hand side square matrix of coefficients and Σ be the 2×2 covariance matrix of the disturbance (which is zero except for the upper

left corner, of course), the unconditional variance matarix of the vector $[y_t, y_{t-1}]'$ is given by

$$\sum_{s=0}^{\infty} B^s \Sigma (B')^s .$$

This can be calculated efficiently by the recursion

$$A_0 = B \qquad\qquad S_0 = \Sigma$$
$$A_{j+1} = A_j^2 \qquad\qquad S_{j+1} = A_j S_j A_j' + S_j .$$

This calculation leads to an unconditional variance for $\tilde{y}$ of 1.4203. Since this is in log units, it is a huge variance, implying a standard deviation of over 100%. However, with the initial value of the trend line in the neighborhood of -900, while the initial value of the data is in the neighborhood of 7, it is clear that the initial conditions are in the extreme tails of the unconditional distribution of $\tilde{y}$. To be more precise, the value for the χ-squared statistic ($z'\Sigma^{-1}z$) for the pair of initial conditions is 708522. Note, though, that the roots of the autoregressive operator are .9999212 and .3713. The half life of the slower decaying component is 8796 quarters, or about 2200 years. So movement toward the trend line is negligible over the sample. The decay from the unusual initial conditions is mainly modeling near-linear trend.

(3) Here's a model for a stock price with dividend δ_t and an i.i.d. "pricing error" or "measurement error" ν_t:

$$P_t = \beta E_t [P_{t+1} + \delta_{t+1}] + \nu_t \qquad (8)$$
$$\delta_t = \delta_{t-1} + \varepsilon_t \qquad (9)$$

The dividend shock ε_t is, like the pricing error ν_t, i.i.d. with mean zero. Assume $\beta \in (0,1)$.

(a) Show that P and δ are cointegrated.

(b) Display the values of cointegrating vector coefficients.

Solving (8) forward produces

$$P_t = \sum_{s=1}^{\infty} \beta^s E_t \delta_{t+s} + \nu_t .$$

From (9) we can conclude that $E_t \delta_{t+s} = \delta_t$ for all $s > 0$ and therefore that

$$P_t = \frac{\beta}{1 - \beta} \delta_t + \nu_t .$$

This means that $(1 - \beta)P_t - \beta\delta_t$ is stationary. But δ itself is clearly, from (9), a non-stationary unit root process, and since P_t is a linear function of δ + a stationary variable, it also is non-stationary. A vector stochastic process that is non-stationary, but has stationary linear combinations, is a cointegrated process by definition.

(c) Display the VECM form of the model, with coefficients explicit functions of β.

The VECM form is simply a reduced-form autoregression in which all variables enter as differences, except that the stationary linear combinations of variables enter as levels, with a less-than-full-rank coefficient matrix. If we let $z_t = [P_t, \delta_t]'$, the VECM form for this model is

$$\Delta z_t = \begin{bmatrix} \frac{1}{1-\beta} \\ 0 \end{bmatrix} \begin{bmatrix} 1 - \beta & \beta \end{bmatrix} z_{t-1} + \begin{bmatrix} (1 - \beta)v_t + \beta\varepsilon_t \\ \varepsilon_t \end{bmatrix}$$