EXERCISE DUE MONDAY, 10/24

The data for this exercise are available on the course website, either as text files or as R data files. The data are seasonally unadjusted CPI for all urban consumers (CPIAUCNS) and seasonally adjusted CPI for all urban consumers (CPIAUCSL). The `cpi.Rdata` file contains both series, as an R time series object. The text files include headers that describe the series, as well as a date column. You will also use the previous exercise's data on the Fed Funds rate in this exercise.

The seasonal effects in CPI data are relatively small. You may get more interesting conclusions if you use seasonally adjusted and unadjusted unemployment data in place of the CPI. I haven't checked that the exercise will work out nicely with those data, but don't see why it shouldn't. You can find unemployment data (and a lot of other macro data) at http://research.stlouisfed.org/fred2/.

(1) Calculate and plot spectral density estimates for inflation calculated from the seasonal and non-seasonal CPI. Use $(1200 * \log p_t / p_{t-1}$, the annualized percentage rate of inflation, as the data series. The plots should be of the log spectral density, or you will not see much except a spike around 0. Price data are usually thought of as not very seasonal, and indeed there are no official adjusted data before 1947. Do you see any evidence of seasonality in the price data? Plot the original series. Note that the variance of inflation obviously declines. The level of the CPI is low at the start of the CPINS series, so the minimum change in the reported index, 0.1, is much larger in percentage terms at the start. How much of the apparent decline in variance can reasonably be attributed to this source? No formal test is expected here, just some sensible calculations.

For this part of the exercise you can use use the built-in `spectrum` or `spec.pgram` commands in R. There is also an R program, `sde.R` that I wrote that is on the web site. My program will only use a square window (sometimes called a Daniell window) and requires Fourier-transformed data as input, but it will handle cross-spectra as well as spectral densities, which the packaged R program does not. My program might also work as a starting point for your own Matlab code. A standard Matlab installation does not include spectral density programs, but you might find something on the web that does it. Matlab has `fft` and convolution programs. If you use programs you haven't written be sure you understand what they are doing, so you don't misapply them. You will need to consider a reasonable bandwidth to use in smoothing your estimates — wide enough to include a reasonable number of harmonic frequencies, narrow enough that, say 3-6 year business cycle frequencies are not in the same smoothing band with the zero frequency.

Date: October 19, 2005.

Calculate a spectral density for `ffr` and cross-spectral densities of `ffr` with `CPINS` and with `CPISD`. Use these to calculate $\tilde{b}$ where b is from the equation $Y = b * X + \varepsilon$. (That is, take the inverse FT of S_{YX}/S_X. Do this for `ffr` in the role of Y and in the role of X, and paired with NS, then SL, CPI. A regression of inflation on the interest rate might be interpreted as a Fisher equation, if the real rate is constant. A regression of the interest rate on inflation might be interpreted as describing monetary policy reaction to past inflation. But both of these interpretations require that only current and past values of the right-hand-side variable appear in the equation. Do any of your estimates show this one-sidedness? Is there any evidence of bias from seasonal adjustment? Note that when you calculate b from the inverse FT, you will have the values of b for negative s at the far end of the series. That is, if you are using 612 data points, 611-612 will correspond to -2 and -1. Note also that it is convenient to have the same number of data points in both series for these frequency domain regression calculations. Otherwise you will need to interpolate one of the FT'd series. It is also convenient to have the sample size an exact multiple of 12, so that if necessary you can remove mean and/or deterministic seasonals just by setting the FT'd data to zero at the $2\pi j/12$ frequencies. Cut a few observations from the beginning of your sample to get a sample of full years.

(2) Construct an ARMA operator of the form $P(L)/Q(L)$, with P second order and Q of order 2 or three, both one-sided and invertible, and with the property that its Fourier transform, $P(e^{-i\omega})/Q(e^{-i\omega})$ has a single large spike at $\pi/2$. Hint: You may want to give numerator and denominator both complex roots, near one in absolute value and near each other, but not equal. After you have constructed your ARMA operator, plot a simulated draw of 100 observations from a process with your operator as its MA operator. You can do this several ways. One approach is to draw i.i.d. shocks and then simulate $Q(L)y_t = P(L)\varepsilon_t$ forward in time. Another is to calculate the spectral density, inverse FT it to get the acf R_X, populate a covariance matrix with R_X values, Cholesky factor it, and multiply an i.i.d. normal vector by that factor.

(3) The Hodrick-Prescott filter applied to a series X_t, $t = 1, \ldots, T$ delivers a filtered "trend" estimate $\hat{X}_t$ that minimizes

$$\sum_{t=1}^{T}(X_t - \hat{X}_t)^2 + \lambda \sum_{t=3}^{T}(\hat{X}_t - 2\hat{X}_{t-1} + \hat{X}_{t-2})^2 \,.$$

As λ goes to ∞, $\hat{X}$ converges to a straight line, and as $\lambda \to 0$ it converges to the X_t sequence itself. If we are not close to the end points of the sample (i.e., to t of 1 or T), the first-order conditions of this minimization problem imply

$$(1 + 6\lambda - 4\lambda(L + L^{-1}) + \lambda(L^2 + L^{-2})\hat{X}_t = X_t \,.$$

Calculate the FT of the Hodrick-Prescott filter for various values of the weight λ on squared second differences. Do this by ignoring endpoints, so that the filter is characterized by $\hat{X}_t = B(L)X_t$, where $B(L)$ is the inverse of the polynomial in L that appears on the left-hand side above. Determine what values, if any, of λ keep the filter's FT close to one for frequencies between zero and 6 years, close to zero outside that range, for quarterly and for monthly data. (Two different values of λ, if any.)