Conditional Forecasts

Lawrence J. Christiano

September 4, 2019

Introduction

- A message of these notes is that the following, apparently simple, question is in fact ambiguous:

  'what will happen if the interest rate is increased
  by 25 basis points (bps) over the following year?'

- Three interpretations:
  - Will set interest rate higher by 25 bps next year, come what may ('Odyssean policy').
  - Expect a configuration of shocks next year which (given the usual monetary policy) will drive up the interest rate by 25 bps.
  - Expect a sequence of shocks to monetary policy (which the public thinks are drawn randomly) that will cause the interest rate to be higher by 25 bps, regardless of the pattern of non-monetary shocks.
- Each interpretation has a different implication for forecasts.

More Examples

- Suppose you have two sets of variables: y_t and x_t.
- Would like to forecast y_{T+j}, $j = 1, 2, ..., f$, conditional on specified future values of at least a subset of x_{T+j}, $j = 1, 2, ..., f$.
- Example 1:
  - x_t denotes the foreign variables in a small open economy model used in the monetary policy division's model and y_t denotes the domestic variables.
  - The analyst has been given forecasts for at least a subset of elements of x_{T+j}, $j = 1, .., f$, by another division and has been asked to 'determine the distribution of future y_t under the assumption that what the other division forecasts actually happens'.
- Example 2: There is concern that oil prices, x_t, will rise in the future and you want to know what that implies for the other variables, y_t.
- Example 3: what happens if there is a shock to foreigners' appetite for domestic versus foreign assets?

Simple Example

- Suppose x_t and y_t are scalars, with the following representation:

$$y_t = a\mu_t + b\nu_t$$
$$x_t = \mu_t + \nu_t,$$

where μ_t and ν_t are iid over time and:

$$\mu_t \sim N\left(0, \sigma_\mu^2\right), \quad \nu_t \sim N\left(0, \sigma_\nu^2\right),$$

$$E\nu_t\mu_{t-j} = 0, \quad j \in (-\infty, \infty).$$

- Suppose that somehow, we've figured out the model parameters:

$$\sigma_\mu, \sigma_\nu, a, b.$$

Simple Example

- At date T we are asked to construct the distribution of Y,

$$Y = \begin{bmatrix} y_{T+1} \\ \vdots \\ y_{T+f} \end{bmatrix},$$

subject to a specific sequence of values of x_t :

$$X = \begin{bmatrix} x_{T+1} \\ \vdots \\ x_{T+f} \end{bmatrix}.$$

- We expect that Y has a multivariate Normal distribution with mean a function of X and variance a function of the model parameters:

$$\sigma_\mu, \sigma_\nu, a, b.$$

Simple Example: Signal Extraction

- The example is sufficiently simple that the distribution of y_t, $t > T$ is simply a function of x_t (no dynamics). For example,

$$E\left[Y|X\right] = \left[\begin{array}{c} E\left[y_{T+1}|x_{T+1}\right] \\ \vdots \\ E\left[y_{T+f}|x_{T+f}\right] \end{array}\right]$$

  – In general the problem is more complicated and the distribution of y_t, $t > T$, is a function of the whole of X, as well as past data.
  – Still, the intuition from the example is sufficient to interpret results from Dynare (see below).

- Conditional on x_t, restriction across shocks:

$$x_t = \mu_t + \nu_t.$$

- Drawing an inference from x_t about the density of μ_t or ν_t is called a *signal extraction problem*.

Conditional Expectation of Y

- We have:

$$E[y_t|x_t] = aE[\mu_t|x_t] + bE[\nu_t|x_t]$$
$$= [a\beta + b(1-\beta)]\, x_t$$
$$\beta = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\nu^2}.$$

- Basic idea: knowledge of x_t shrinks set of shocks, μ_t, ν_t, that can occur.
  - Given x_t, the most likely value of the shocks is (using the Normality assumption): ▸ details

$$\hat{\mu}_t = \text{argmax}_{\mu_t} f(\mu_t|x_t) = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\nu^2} x_t, \qquad \hat{\nu}_t = \frac{\sigma_\nu^2}{\sigma_\mu^2 + \sigma_\nu^2} x_t$$

  - Standard signal extraction.

Getting a Conditional Expectation in Dynare

- Setup:
  - put NaN (MATLAB for 'not a number') in each entry of Y.
  - put numbers in the elements of X corresponding to the conditioning information that you have. If there are elements of x_t that you know nothing about, put in NaN.
- Provide the historical data on x_t and y_t and the 'data', Y and X.
  - You now have one big data set with missing observations.
  - Implement Dynare's 'estimate' command.
  - Provide all your data to Dynare in the usual way.
  - Set mode_compute=0 and mh_replic=0, but include the 'smoother' command.
  - The structure, oo_.SmoothedVariables, contains the results Dynare's Kalman smoother.
    - All entries with NaN will be replaced by estimates.

Experiments to Illustrate the Results

- Take our simple NK model (DS version), with an iid monetary policy shock, u_t.
- Assign parameter values:

  $\beta = 0.99, \phi_x = 0, \phi_\pi = 1.5, \alpha = 0.85, \rho = 0.9, \lambda = 0.7, \varphi = 1, \theta$
  $\sigma_a = \sigma_\tau = \sigma_u = 0.01.$

Odyssean Policy versus Smoothing

Figure: Baseline Parameterization

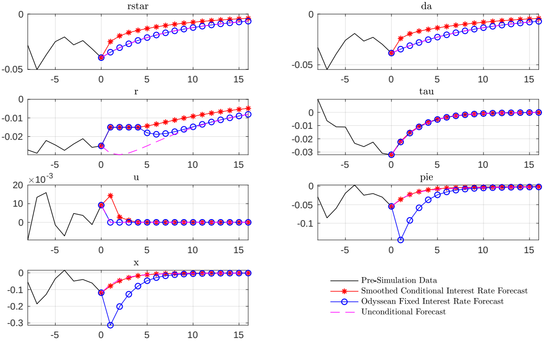

Note: with the Odyssean policy and unconditional forecast, all future shocks are set to zero. With the conditional expectation, the procedure picks the most likely configuration of shocks, that explain the fixed interest rate. Note the Kalman smoother does not employ the tau shock at all. The most important shock for r is the ε_t^a shock, and that's the one that the smoother puts the most weight on.

Odyssean Policy versus Smoothing

Figure: Only the Monetary Policy Shock Matters

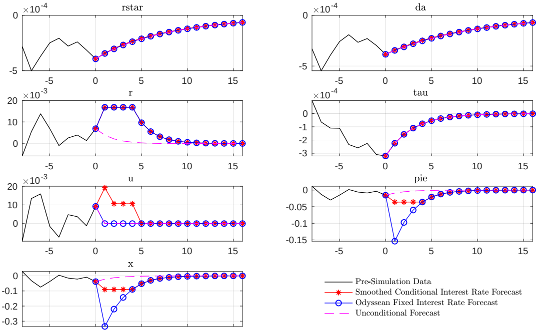

Note: Here, $\sigma_a = \sigma_\tau = 0.0001$, and $\sigma_u = 0.01$. This has no impact on unconditional forecast and Odyssean policy. There is a substantial impact on the conditional expectation. Given the specification of shock variances, Dynare has been 'tricked' into thinking the monetary policy shock is the only important shock. So, it uses the monetary policy shock to get the interest rate to follow its path. The impact of the Odyssean policy is bigger on output because in that case the monetary authority commits to sticking to the interest rate path, while in the conditional expectations, people are fooled in each period to think that the interest rate is high only temporarily (the monetary policy shock is - according to their belief - only temporary).

Conclusion

–

- We've described three ways to interpret the question, 'what will happen if the interest rate remains fixed over the next year or two?"
  - Also, have addressed a large class of conditional forecasting questions.
- The Dynare code, smallNK_model.mod implements all three.
  - It is contained in the file, forecasting.zip.

Simple Example: Bayes' Rule

- Let $p(x)$, $p(\mu)$ denote the marginal density of x_t and μ_t respectively (from here on, we drop the t subscript).
- Let $p(x|\mu)$ denote the distribution of x conditional on observing μ.
- Our assumptions imply:

$$p(x|\mu) = \frac{1}{\sqrt{2\pi\sigma_\nu^2}} exp\left\{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma_\nu^2}\right\}$$

$$p(\mu) = \frac{1}{\sqrt{2\pi\sigma_\mu^2}} exp\left\{-\frac{1}{2}\frac{\mu^2}{\sigma_\mu^2}\right\}$$

$$p(x) = \frac{1}{\sqrt{2\pi\left(\sigma_\mu^2+\sigma_\nu^2\right)}} exp\left\{-\frac{1}{2}\frac{x^2}{\sigma_\mu^2+\sigma_\nu^2}\right\}.$$

Simple Example: Bayes' Rule

- With a little algebra, the marginal density of μ given x :

$$p(\mu|x) = \frac{p(x|\mu)\,p(\mu)}{p(x)} = \frac{1}{\sqrt{2\pi s^2}} exp\left\{ -\frac{1}{2}\frac{\left(\mu - \frac{\sigma_\mu^2}{\sigma_\mu^2+\sigma_\nu^2}x\right)^2}{s^2} \right\},$$

where βx is the value of μ that maximizes $p(\mu|x)$ and

$$s^2 = \frac{\sigma_\mu^2\sigma_\nu^2}{\sigma_\mu^2+\sigma_\nu^2}, \quad E[\mu|x] = \beta x, \quad \beta = \frac{\sigma_\mu^2}{\sigma_\mu^2+\sigma_\nu^2}.$$

- Similarly,

$$E[\nu|x] = (1-\beta)\,x$$

Go Back