

THE MINNESOTA PRIOR

1. THE MINNESOTA PRIOR: NOTATION

$$y(t) = c + \sum_{s=1}^k B_s y(t-s) + \varepsilon(t)$$

with data for $t = 1, \dots, T$ and $\varepsilon(t) \mid \{y(t-s), s \geq 1; B, c, \Sigma\} \sim N(0, \Sigma)$.

2. MN PRIOR: SINGLE-UNIT-ROOT PRIOR

We introduce data for the artificial date t^* in which

$$y(t^*) = y(t^* - 1) = \dots = y(t^* - k) = \underset{n \times 1}{\bar{y}} \lambda \underset{n \times 1}{\lambda}$$

and the vector of 1's that corresponds to the constant term in the data matrix is set to λ in the t^* observation. The vector \bar{y} is usually set to the sample mean of the initial conditions, i.e.

$$\bar{y} = \frac{1}{k} \sum_{s=1}^k y(1-s).$$

3. REGRESSION EQUATION FORM

$$\begin{aligned} \bar{y}\lambda &= \left(\sum_{s=1}^k B_s \right) \bar{y}\lambda + c\lambda + \varepsilon(t^*) \\ (I - B(1))\bar{y} &= c + \varepsilon(t^*). \end{aligned}$$

Thus the prior is centered on a part of the parameter space where either $c = 0$ and the system contains a unit root with \bar{y} as its eigenvector, or $c \neq 0$, y is stationary, and $y(0)$ is close to the model's implied population mean.

4. CONNECTING BELIEFS ABOUT c AND B

One may want to favor the unit-root part of this region. Can do so by omitting the "constant" from the dummy observation, by adding a dummy observation that asserts a prior directly on c , by adding a dummy observation that omits the constant in addition to the one that includes it, or by including the multiple-unit-root dummy observations.

It is a good property of the single-unit-root dummy that it makes beliefs about c connect to beliefs about the presence of non-stationarity.

Date: November 28, 2006.

©2006 by Christopher A. Sims. This document may be reproduced for educational and research purposes, so long as the copies contain this notice and are retained for personal use or distributed free.

Unit roots convert constants into polynomial trends.

5. MN PRIOR: MULTIPLE UNIT ROOTS

$$y(t_j^*) = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \bar{y}_j \lambda \\ 0 \\ \vdots \\ 0 \end{bmatrix} = y(t_j^* - 1) = \dots = y(t_j^* - k),$$

with the \bar{y}_j in the bracketed vector occurring in the j 'th position, and the column of the data matrix corresponding to c set to 0.

6. REGRESSION EQUATION FORM

This j 'th dummy observation reads, for the j 'th equation,

$$\lambda \bar{y}_j = \sum_s B_{jjs} \bar{y}_j \lambda + \varepsilon_j(t_j^*)$$

and for equation $k \neq j$,

$$0 = \sum_s B_{kjs} \bar{y}_j + \varepsilon_k(t_j^*).$$

It is easy to see that this prior favors B 's whose off-diagonal elements are small, and also B 's that are close to putting a unit root in the $B_{jj}(L)$ polynomial for each j .

7. MN PRIOR: THE ORIGINAL

The original Minnesota Prior postulates a separate set of dummy observations for each equation, implementing a prior that specifies independent distributions for all coefficients, with the coefficient B_{jks} having mean 0 unless $j = k$ and $s = 1$, in which case the mean is 1, and variance $\pi_1^{-1} \pi_2^{1-\delta(j,k)} s^{-\pi_3}$, where $\delta(j, k)$ is a function equal to 1 if $j = k$ and 0 otherwise. With $\pi_2 < 1$, this prior specifies that coefficients on own lags ($B_{jj}(s)$, $s = 1, \dots, k$) are likely to be larger in absolute value than coefficients on other variables, and with $\pi_3 > 0$, it specifies that coefficients on more distant lags are likely to be smaller.

The difference in dummy observations across equations implies that equation-by-equation estimation is not ML, and this fact was ignored in the early literature. If we omit the π_2 term (that is, set it to 1), then own lag coefficients and cross-variable coefficients are treated symmetrically, and the prior can be implemented with system-wide dummy observations.

8. SYSTEM DUMMY OBSERVATIONS FOR $\pi_2 = 1$ MN PRIOR

If t_{ks}^* is the artificial date for the dummy observation applying to the s 'th lag of the k 'th variable, one of these dummy observations has the form

$$(1) \quad y(t_{k1}^*) = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \pi_1 \sigma_k \\ 0 \\ \vdots \\ 0 \end{bmatrix} = y(t_{k1}^* - 1), \quad y(t_{ks}^* - s) = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \pi_1 s^{\pi_3} \sigma_k \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

$$(2) \quad y(t_{ks}^* - v) = 0 \text{ for } v \neq s > 1 \text{ and for } v > s = 1.$$

9. SETTING σ_k

The parameter σ_k is a measure of the degree of variability in the k 'th variable. It has most commonly, again somewhat inconsistently, been set by estimating a low-order univariate AR regression for each variable and taking σ_k to be the standard deviation of the residual of that regression. It could also be set as the sample standard deviation of the initial conditions for each variable, or on the basis of rough a priori reasoning about the likely scale of variation in each variable. Usually only the order of magnitude matters.