

Structural dynamic factor models

Tyler Sotomayor

2024

These solutions calculate the covariance lost under rank reduction and the units of a normalized policy response. A scalar filtering example follows a missing measurement through the prediction step; a final calculation applies the factor-count penalty to three candidate models.

Based on the 2024 BSE course taught by Luca Gambetti and Konstantin Boss. I have revised the classroom code and added the accompanying explanations. The notes identify the papers, data, and assumptions used in each application. These are my course notes, not an official BSE publication. A classroom replication should not be read as a replication of every result in its underlying paper.

1 What rank reduction discards

The retained impact matrix is

$$R = \begin{pmatrix} 3 & 0 \\ 0 & 2 \\ 0 & 0 \end{pmatrix}, \quad RR' = \text{diag}(9, 4, 0).$$

The discarded covariance is $\text{diag}(0, 0, 1)$. The retained trace share is $(9 + 4)/(9 + 4 + 1) = 13/14 \simeq 0.928571$. It is incorrect to write $RR' = \Sigma_F$ because the third eigenvalue is positive.

With $K_q = (e_1, e_2)$ and $M_q = \text{diag}(3, 2)$, the retained standardized shocks for $(3, 2, 5)'$ are $w = M_q^{-1}K_q'\varepsilon = (1, 1)'$. Reconstruction gives $Rw = (3, 2, 0)'$, leaving discarded innovation $(0, 0, 5)'$. A low-rank approximation can explain a large fraction of unconditional covariance while missing a large component of a particular realized observation.

2 Policy-shock units

The reporting multiplier is $0.5/0.2 = 2.5$. At horizon four, the GDP response becomes $2.5(-0.003) = -0.0075$ log points, or approximately -0.75 percent. The exact proportional change would be $100(e^{-0.0075} - 1) \simeq -0.7472$ percent.

The original shock series still has unit variance. If the impulse response is rescaled, its innovation-unit interpretation changes and the multiplier must be recorded. The other shock columns describe different disturbances, so applying the policy multiplier to them does not standardize them to an economically comparable intervention. The replication retains unit-variance shocks and stores `policyScale` separately.

3 A missing observation

At date one, the innovation is one and $F_1 = 1 + 0.25 = 1.25$. Thus $K_1 = 1/1.25 = 0.8$, the posterior mean is 0.8, and its variance is $(1 - 0.8)1 = 0.2$. The Gaussian log-likelihood contribution is

$$\ell_1 = -\frac{1}{2}[\log(2\pi) + \log(1.25) + 1/1.25] \simeq -1.430510.$$

At date two, prediction gives mean $0.8(0.8) = 0.64$ and variance $0.8^2(0.2) + 0.36 = 0.488$. No measurement is available, so these are also the posterior moments and the measurement log-likelihood contribution is zero.

At date three, the predicted mean is $0.8(0.64) = 0.512$ and variance is $0.8^2(0.488) + 0.36 = 0.67232$. Now

$$K_3 = \frac{0.67232}{0.67232 + 0.25} \simeq 0.728944,$$

so the zero observation yields posterior mean $0.512(1 - K_3) \simeq 0.138780$ and variance $0.67232(1 - K_3) \simeq 0.182236$. Deleting the missing second row would apply only one transition between the first and third measurements and produce a different, incorrectly timed answer.

4 The factor penalty

The penalty per factor is

$$\frac{100 + 200}{100 \cdot 200} \log(100) = 0.015 \log(100) \simeq 0.069078.$$

Consequently,

$$IC_2(1) = \log(0.50) + 0.069078 \simeq -0.624070,$$

$$IC_2(2) = \log(0.46) + 2(0.069078) \simeq -0.638373,$$

$$IC_2(3) = \log(0.43) + 3(0.069078) \simeq -0.636737.$$

The minimum is at two factors. The third factor improves reconstruction, but not enough to compensate for its additional penalty. Unpenalized reconstruction error would select three among these candidates and, more generally, would keep falling as the retained PCA rank increases. This is precisely why reconstruction fit alone cannot determine a parsimonious factor dimension.