

Quantile SVARs and growth risk

Tyler Sotomayor

2024

The exercises derive the quantile check loss and its dual bound, distinguish a width innovation from its impact vector, impose contemporaneous zeros, and check the dates of the growth forecast. Each solution shows the intermediate calculation and its interpretation.

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 The forecast origin

There are $4(2019 - 1960) + 4 = 240$ predictor quarters. To use both t and $t - 1$, the first origin is 1960 Q2. The last remains 2019 Q4, giving 239 observations. Adding one quarter to each origin produces targets from 1960 Q3 to 2020 Q1. Growth at the last target is $100[\log GDP_{2020Q1} - \log GDP_{2019Q4}]$, so the first-quarter 2020 GDP level is needed even though it is never a current predictor.

The count can be checked directly from the target calendar: $4(2020 - 1960) + (1 - 3) + 1 = 239$. Plotting the origins as if they were target dates would shift the fitted growth series back by one quarter.

2 A quantile certificate

The sample median is $b = 3$, with residuals $(-2, -1, 0, 1, 6)'$. At $\tau = 1/2$, the objective is

$$P = \frac{1}{2}(2 + 1 + 0 + 1 + 6) = 5.$$

Take $a = (-1/2, -1/2, 0, 1/2, 1/2)'$. Every weight is in the allowed box and $\mathbf{1}'a = 0$, which is the dual equality constraint for an intercept. The dual objective is

$$D = y'a = -\frac{1}{2} - 1 + 0 + 2 + \frac{9}{2} = 5.$$

Weak duality gives $D \leq P^* \leq P$ for the unknown optimal primal value P^* . Since both bounds equal five, $P^* = 5$. No differentiability or local curvature argument is needed. The automated test repeats this calculation and checks the returned dual weights and objective gap.

3 A normalized distribution innovation

First, $\gamma\Sigma\gamma' = 4$, so the unrestricted impact vector is $\Sigma\gamma'/2 = (2, 1/2)'$. The scalar shock $u_1/2$ has variance one. Its second-variable impact is not zero because the reduced-form innovations have covariance one.

For $D = (0, 1)$, $D\Sigma D' = 1$ and $\gamma\Sigma D' = 1$. Hence

$$\alpha = (1, 0) - 1 \cdot 1^{-1}(0, 1) = (1, -1).$$

Its variance is $\alpha \Sigma \alpha' = 4 - 1 - 1 + 1 = 3$, while $\Sigma \alpha' = (3, 0)'$. The normalized restricted impact is therefore $(\sqrt{3}, 0)'$. The zero is exact under the restriction, not a small estimated coefficient. Renormalization matters: using the old standard deviation two would no longer produce a unit-variance shock.

4 The covariance denominator

Three lags leave $239 - 3 = 236$ residual observations. There are $1 + 7 \times 3 = 22$ regressors in each equation. The unbiased divisor is $236 - 22 = 214$. The classroom expression instead gives $240 - 3 - 21 - 1 = 215$. For the same residual cross product, $\hat{\Sigma}_{preferred} = (215/214)\hat{\Sigma}_{source}$.

For any fixed γ , replacing Σ by $c\Sigma$ yields

$$\frac{c\Sigma\gamma'}{\sqrt{\gamma c\Sigma\gamma'}} = \sqrt{c} \frac{\Sigma\gamma'}{\sqrt{\gamma\Sigma\gamma'}}.$$

The preferred response is larger by $\sqrt{215/214}$, approximately 1.002334. The projected direction is unchanged because the scalar covariance factor cancels from its regression projection. The benchmark removes this known scale change before comparing the two implementations.